



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-498359

Advanced Simulation and Computing FY12-13 Implementation Plan, Volume 2, Revision 0.5

R. Meisner, J. Phillips, M. McCoy, R. Webster

September 9, 2011

Disclaimer


This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

FY12-FY13 Implementation Plan

Volume 2, Rev. 0.5

September 26, 2011



LLNL-TR-498359

NA-ASC-120R-11-Vol.2-Rev.0.5-IP

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.



Advanced Simulation and Computing

FY12–13 IMPLEMENTATION PLAN

Volume 2, Rev. 0.5

September 26, 2011

Approved by:

Robert Meisner,

NNSA ASC Program Director

Signature

9/23/11

Date

Wendy Cieslak,

SNL ASC Executive

Concurrence Received

9/22/11

Michel McCoy,

LLNL ASC Executive

Concurrence Received

9/22/11

Robert Webster,

LANL ASC Executive

Concurrence Received

9/22/11

ASC Focal Point

Robert Meisner

NA-114

Tele.: 202-586-0908

FAX: 202-586-0405

bob.meisner@nnsa.doe.gov

IP Focal Point

Atinuke Arowojolu

NA-114

Tele.: 202-586-0787

FAX: 202-586-7754

atinuke.arowojolu@hq.doe.gov

Implementation Plan Contents at a Glance

Section No./Title		Vol. 1	Vol. 2
I.	Executive Summary		
II.	Introduction		
III.	Accomplishments		
IV.	Product Descriptions		
V.	ASC Level 1 and 2 Milestones		
VI.	ASC Roadmap Drivers for FY11–FY12		
VII.	ASC Performance-Based Initiatives/Multi-Site Target Drivers for FY12–FY13		
VIII.	Budget		
Appendix A.	Glossary		
Appendix B.	Codes		
Appendix C.	Points of Contact		
Appendix D.	Academic Alliance Centers		
Appendix E.	ASC Obligation/Cost Plan		

Contents

I. EXECUTIVE SUMMARY	1
II. INTRODUCTION	2
ASC Contributions to the Stockpile Stewardship Program.....	3
III. ACCOMPLISHMENTS FOR FY10–FY11.....	5
Computational Systems and Software Environment	5
Facility Operations and User Support	7
Academic Alliances	8
IV. PRODUCT DESCRIPTIONS BY THE NATIONAL WORK BREAKDOWN STRUCTURE	11
<i>WBS 1.5.4: Computational Systems and Software Environment.....</i>	<i>11</i>
WBS 1.5.4.1: Capability Systems	11
WBS 1.5.4.1 Systems Requirements and Planning (LANL)	11
WBS 1.5.4.1 Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (SNL, LANL).....	12
WBS 1.5.4.1 Alliance for Computing at Extreme Scale Architecture Office (LANL, SNL).....	13
WBS 1.5.4.2: Capacity Systems	15
WBS 1.5.4.2 Production Planning and Integration (LLNL)	15
WBS 1.5.4.2 Production Planning and Integration (LANL).....	16
WBS 1.5.4.2 ASC Capacity Systems (SNL).....	17
WBS 1.5.4.3: Advanced Systems	18
WBS 1.5.4.3 Petascale Application Enablement (LLNL)	18
WBS 1.5.4.3 BlueGene/P and BlueGene/Q Research and Development (LLNL)	19
WBS 1.5.4.3 Hyperion Test Bed (LLNL).....	20
WBS 1.5.4.3 Advanced Systems Technology Research and Development (SNL)	21
WBS 1.5.4.3 Heterogeneous Computing (SNL).....	22
WBS 1.5.4.4: System Software and Tools	23
WBS 1.5.4.4 System Software Environment for Scalable Systems (LLNL).....	23
WBS 1.5.4.4 Applications Development Environment and Performance Team (LLNL)	24
WBS 1.5.4.4 Co-Design through Mini-Applications (LANL).....	25
WBS 1.5.4.4 Programming Models for the Next-Generation Scientific-Computing Environment (LANL).....	27
WBS 1.5.4.4 High Performance Computing Systems Research (LANL)	28
WBS 1.5.4.4 Application Readiness (LANL).....	29
WBS 1.5.4.4 Software Support (LANL).....	30
WBS 1.5.4.4 Software and Tools for Scalability and Performance (SNL).....	31
WBS 1.5.4.4 Resilience (SNL)	32
WBS 1.5.4.4 System Simulation and Computer Science (SNL)	33
WBS 1.5.4.4 Application Performance Analysis with Mantevo Proxy Applications (SNL)	35
WBS 1.5.4.4 Scalable, Fault-Resilient Programming Models (SNL).....	36

WBS 1.5.4.5: Input/Output, Storage Systems, and Networking	37
WBS 1.5.4.5 Archive Storage (LLNL)	37
WBS 1.5.4.5 Parallel and Network File Systems (LLNL)	38
WBS 1.5.4.5 Networking and Test Beds (LLNL)	39
WBS 1.5.4.5 Archival and File Systems (LANL)	40
WBS 1.5.4.5 Archival Storage (SNL)	42
WBS 1.5.4.5 Scalable Input/Output Research (SNL)	42
WBS 1.5.4.5 Scalable Interconnects for Extreme-Scale Tightly Coupled Systems (SNL)	44
WBS 1.5.4.6: Post-Processing Environments	45
WBS 1.5.4.6 Scientific Visualization (LLNL)	45
WBS 1.5.4.6 Scientific Workflow and Data Management (LLNL)	46
WBS 1.5.4.6 Production Visualization (LANL)	47
WBS 1.5.4.6 Visualization and Data Analysis Research and Development Project (LANL)	49
WBS 1.5.4.6 Scalable Data Analysis (SNL)	50
WBS 1.5.4.7: Common Computing Environment (Moved)	51
WBS 1.5.5: Facility Operations and User Support	52
WBS 1.5.5.1: Facilities, Operations, and Communications (Retired)	52
WBS 1.5.5.2: User Support Services	52
WBS 1.5.5.2 Hotlines and System Support (LLNL)	52
WBS 1.5.5.2 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)	53
WBS 1.5.5.2: User Support (SNL)	54
WBS 1.5.5.3: Collaborations	56
WBS 1.5.5.3 Program Support (LLNL)	56
WBS 1.5.5.3 Program Support (LANL)	57
WBS 1.5.5.3 Program Support (SNL)	57
WBS 1.5.5.3 Applications in Support of Manufacturing Production and Connectivity (Y-12)	58
WBS 1.5.5.4: System and Environment Administration and Operations	60
WBS 1.5.5.4 System and Environment Administration and Operations (LLNL)	60
WBS 1.5.5.4 System Administration and Storage (LANL)	61
WBS 1.5.5.4 Operations and Procurement Support (LANL)	62
WBS 1.5.5.4 Requirements Planning (LANL)	63
WBS 1.5.5.4 Computing Platform Integration and Deployment (LANL)	63
WBS 1.5.5.4 Production Computing Services (SNL)	64
WBS 1.5.5.5: Facilities, Network, and Power	66
WBS 1.5.5.5 Facilities, Network, and Power (LLNL)	66
WBS 1.5.5.5 Facilities, Networking, and Power (LANL)	67
WBS 1.5.5.5 Facilities, Networking, and Power (SNL)	69
WBS 1.5.5.6: Common Computing Environment	70
WBS 1.5.5.6 Tripod Operating System Software	70
WBS 1.5.5.6 Open/SpeedShop	71
WBS 1.5.5.6 Workload Characterization	72
WBS 1.5.5.6 SARAPE Redevelopment	73
WBS 1.5.5.6 Distributed Data Services	74
WBS 1.5.5.6 Shared Workspace	75
WBS 1.5.5.6 Debugger	75
WBS 1.5.5.6 Inter-Site High Performance Computing	76
WBS 1.5.5.6 Open Source Contract Maintenance	78

WBS 1.5.5.6 Open Message Passing Interface Integration/Scaling	79
WBS 1.5.5.6 File System Architecture and Integration	80
WBS 1.5.5.6 General-Purpose Computing on Graphics Processing Units Programming Environment	81
WBS 1.5.5.6 Resource Manager Capability Integration	81
V. ASC LEVEL 1 AND 2 MILESTONES.....	83
VI. ASC PERFORMANCE-BASED INITIATIVES/MULTI-SITE TARGET DRIVERS FOR FY12–FY13	110
VII. PERFORMANCE MEASURES	111
VIII. BUDGET	112
APPENDIX A. GLOSSARY.....	116
APPENDIX C. POINTS OF CONTACT	120
APPENDIX D. WBS 1.5.1.4-TRI-001 ACADEMIC ALLIANCE CENTERS	123
California Institute of Technology	123
Purdue.....	124
Stanford University	126
University of Michigan	127
University of Texas	128
APPENDIX E. ASC OBLIGATION/SPEND PLAN	130

I. Executive Summary

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the surety and reliability of the U.S. nuclear stockpile. The SSP uses nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering (D&E) programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program requires the continued use of experimental facilities and programs, and the computational enhancements to support these programs.

The Advanced Simulation and Computing Program (ASC) is a cornerstone of the SSP, providing simulation capabilities and computational resources that support annual stockpile assessment and certification, study advanced nuclear weapons design and manufacturing processes, analyze accident scenarios and weapons aging, and provide the tools to enable stockpile Life Extension Programs (LEPs) and the resolution of Significant Finding Investigations (SFIs). This requires a balanced resource, including technical staff, hardware, simulation software, and computer science solutions.

In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. In its second decade, ASC is now focused on increasing predictive capabilities in a three-dimensional (3D) simulation environment while maintaining support to the SSP. The program continues to improve its unique tools for solving progressively more difficult stockpile problems (sufficient resolution, dimensionality, and scientific details), quantify critical margins and uncertainties, and resolve increasingly difficult analyses needed for the SSP.

Moreover, ASC's business model is integrated and focused on requirements-driven products that address long-standing technical questions related to enhanced predictive capability in the simulation tools.

ASC must continue to meet three objectives:

- **Objective 1. Robust Tools.** Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs, LEPs, annual assessments, and evolving future requirements.
- **Objective 2. Prediction through Simulation.** Deliver verified and validated physics and engineering tools to 1) enable simulations of nuclear weapons performance in a variety of operational environments and physical regimes, and 2) enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
- **Objective 3. Balanced Operational Infrastructure.** Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work (DSW) and SSP needs for production and advanced/high-end simulation capabilities.

II. Introduction

Prior to the start of the nuclear testing moratorium in October 1992, the nuclear weapons stockpile was maintained through 1) underground nuclear testing and surveillance activities and 2) “modernization” (in other words, development of new weapons systems). A consequence of the nuclear test ban is that the safety, performance, and reliability of U.S. nuclear weapons must be ensured by other means for systems far beyond the lifetimes originally envisioned when the weapons were designed.

Following the start of the moratorium, the ASC Program now supports the National Nuclear Security Administration’s (NNSA’s) overarching goal of nuclear weapons stewardship: *“We continue to advance the Stockpile Stewardship Program to push the scientific and engineering boundaries needed to maintain our nuclear arsenal. It also means maintaining the basic science and engineering that is the foundation of the weapons program.”*¹

In 1996, ASCI—the Accelerated Strategic Computing Initiative—was established as an essential element of the SSP to provide nuclear weapons simulation and modeling capabilities.

In 2000, the NNSA was established to carry out the national security responsibilities of the Department of Energy (DOE), including maintenance of a safe, secure, and reliable stockpile of nuclear weapons and associated materials capabilities and technologies.

Shortly thereafter, in 2002, ASCI matured from an initiative to a recognized program and was renamed the ASC Program.

NNSA will carry out its responsibilities through the twenty-first century in accordance with the current Administration’s vision and the Nuclear Posture Review (NPR) guidance. NNSA Administrator Thomas P. D’Agostino summarized² the NNSA objectives for the SSP as follows:

“Our fundamental national security responsibilities for the United States include:

- *Assuring the safety, security and reliability of the U.S. nuclear weapons stockpile while at the same time transforming the stockpile and the infrastructure that supports it;*
- *Reducing the threat posed by nuclear proliferation; and,*
- *Providing reliable and safe nuclear reactor propulsion systems for the U.S. Navy.”*

“Throughout the past decade, the Stockpile Stewardship Program has proven its ability to successfully sustain the safety, security and reliability of the nuclear arsenal without

¹ *NNSA Strategic Planning Guidance for FY2010–2014*, April 2008, page 17.

² Testimony on the *FY 2008 National Defense Authorization Budget Request for the Department of Energy’s NNSA* before the House Armed Services Subcommittee, March 20, 2007.

resorting to underground nuclear testing. The SSP also enables the U.S. to provide a credible strategic deterrent capability with a stockpile that is significantly smaller.”

Additionally, key investment recommendations cited in the 2010 NPR are to:

- Strengthen the science, technology, and engineering base needed for conducting weapon system LEPs
- Mature advanced technologies to increase weapons surety
- Qualify weapon components and certify weapons without nuclear testing
- Provide annual stockpile assessments through weapons surveillance

This strategy includes developing and sustaining high-quality scientific staff, as well as supporting computational and experimental capabilities.³

The ASC Program plays a vital role in the NNSA infrastructure and its ability to respond to the NPR guidance. The program focuses on developing modern simulation tools that can provide insights into stockpile problems, providing tools with which designers and analysts can certify nuclear weapons, and guiding any necessary modifications in nuclear warheads and the underpinning manufacturing processes. Additionally, ASC is enhancing the predictive capability necessary to evaluate weapons effects, designing experiments, and ensuring test readiness.

ASC continues to improve its unique tools to solve progressively more difficult stockpile problems—with a focus on sufficient resolution, dimensionality, and scientific details—to enable Quantification of Margins and Uncertainties (QMU) and to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW program provides requirements for simulation. These requirements include planned LEPs, stockpile support activities and requirements for future capabilities. ASC’s advanced, leading-edge technologies in high performance computing (HPC) and predictive capability meet DSW’s short- and long-term needs, including the annual assessments and certifications, as well as SFIs.

The following section provides an overview of FY10 and upcoming ASC contributions to the SSP:⁴

ASC Contributions to the Stockpile Stewardship Program

In FY11, ASC continued delivering science-based simulation tools for annual assessments and next-generation LEPs, focusing on improved physics, fidelity, and calculations in support of DSW and the National Code Strategy. The methodology for predictive capability assessment was demonstrated in FY11 for a limited set of simulations common to both physics laboratories. ASC assessed the ability to simulate full system performance near thresholds where data are sparse. Cielo, the capability HPC system, replaced Purple and began operation. Cielo was upgraded to 1.37 petaFLOP/s,

³ 2010 Nuclear Posture Review Report, April 2010, p. 42.

⁴ FY11 ASC Program Plan for past, present, and planned contributions to the SSP.

the system was accepted, and the second round of the Capability Computing Campaign (CCC) was initiated. Sequoia, the advanced architecture system, received a “go” from the program’s “go/no go” review. Contract negotiations and delivery schedule have been finalized. As of September 30, 2011, facility site preparation work was well underway, and design and procurement of the parallel file system was complete. Building on the FY10 Scalable Applications Project (SAP) milestone, the SAP effort has extended the knowledge base, documentation, and training to enable tri-lab code teams to use Sequoia. Successful procurement of the next-generation tri-lab Linux capacity clusters (TLCC2) and the associated common user environment milestones were executed across the three NNSA laboratories. Initial investments were made towards the proposed joint DOE Office of Science and NNSA exascale initiative. Proposal responses to ASC’s Request for Information were evaluated in preparation for the next round of Predictive Science Academic Alliance Program (PSAAP) engagements.

In FY12, ASC intends to release a Request for Proposal and conduct selection activities towards the PSAAP II program. Cielo, the new capability machine in operations at LANL, will complete the second round of CCC and is expected to receive approval of the project closeout (CD4) documentation. Investments in exascale-critical technologies are expected to continue. ASC will continue to provide a reliable, available, and secure environment for distance computing. On the Physics and Engineering Models (PEM) front, a Level 1 milestone to advance capabilities for annual assessments and resolution of SFIs associated with early-phase primary implementation will continue. Verification and validation assessment of improvements in primary performance codes for boost will continue, and the subprogram will continue to strategize and complete common modeling to validate improvements in support of the National Boost Initiative. The Integrated Codes (IC) program element will concentrate on preparations for the next predictive capability framework (PCF) pegpost and on a Level 1 milestone for simulating late-time primary implosion and initial explosion, as well as on scalability enhancements targeting future computing platforms. Staged delivery of all Sequoia racks is expected by the end of the third quarter. The system is projected to be integrated on the unclassified network by the end of the fourth quarter. Integration and operation of TLCC2 systems will begin in FY12.

In FY13 and beyond, ASC will be focusing on strengthening the science basis and driving down uncertainties for weapons simulations to a degree that NNSA can ultimately, and credibly, claim predictive capability; instituting a robust, formalized peer review system; increasing the number of production computing cycles to support increased use of simulation in stockpile activities and reliance on uncertainty quantification (UQ) in weapons decisions; and pursuing next-generation computing to meet time-urgent, future predictive science capability needs as documented in the *ASC Platform Strategy* and the *Predictive Capability Framework*.

III. Accomplishments for FY10–FY11

ASC accomplishments from quarter 4, fiscal year 2010, through quarter 3, fiscal year 2011, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Performed advanced applications preparation work for Sequoia, including particular focus on threading, transactional memory (TM), and methods to exploit short vector floating point unit
- Planned and executed a tri-lab procurement process for the next capacity tri-lab hardware platform; supported the systems software and tools on the current systems
- Supported ASC exascale planning, including organizing an exascale workshop in San Francisco and leading tri-lab working group efforts in tools and in applications; the workshop white papers are available on the Web⁵
- Received an R&D 100 Award for the Stack Trace Analysis Tool (STAT), a highly scalable debugging technology that identifies process behavioral equivalence classes (received in collaboration with researchers at two universities)

LANL Accomplishments for Computational Systems and Software Environment

- Performed activities critical to Cielo deployment as a tri-lab user facility capability machine, including procurement, delivery, acceptance testing, and application readiness support
- Supported the exascale planning effort, leading tri-lab working groups in programming models and in visualization and data analysis to produce white papers and presentation for the ASC community
- Developed mini-applications (mini-apps) as part of a new co-design through mini-apps effort for use as vehicles for exploring the challenges of next-generation computing technologies
- Developed and applied enhanced program analysis and transformation techniques as new approaches to understanding and managing the complex memory hierarchy of many-core processors
- Delivered visualization infrastructure enhancements, including the upgraded visualization cluster, Viewmaster 2, and visualization capabilities at distances not

⁵ <https://asc.llnl.gov/exascale/papers.php>

before reached (from the Strategic Computing Complex (SCC) to the National Security Sciences Building (NSSB))

- Delivered advanced data structures and parallelization techniques to the Eulerian Applications project,

SNL Accomplishments for Computational Systems and Software Environment

- Received an NNSA Environmental Stewardship Award for prior-year's power-related research and development (R&D)
- Collaborated to interactively explore the largest-ever adaptive mesh refinement (AMR) visualization using the CTH code, which consisted of 3 billion cells and 32 thousand files—data that was generated from a simulation run on Cielo
- Determined with runs up to 64K cores on Cielo that six ASC applications from the tri-labs delivered an average performance improvement of 9.6 times the prior capability platform
- Developed and executed new mini-app, miniGhost, on over 100,000 cores as part of Cielo phase 2 acceptance testing
- Expanded the capabilities of the Structural Simulation Toolkit (SST) suite with improved network congestion models; enhanced message passing interface (MPI) tracing with DUMPI, a new M5 processor simulator, more memory models, and several system technology models to analyze reliability, thermal, and energy usage

Tri-Lab Accomplishments for Computational Systems and Software Environment

- Successfully completed the Level 2 milestone 3954, *Deployment of a common capacity computing environment*
- Released Tripod Operating System Software (TOSS) 1.4 (based on Red Hat Enterprise Linux (RHEL) 5.5) production version for general availability (GA)
- Deployed Open|SpeedShop (O|SS) on all tri-lab production platforms and provided beta release of the framework targeted for scalability
- Deployed the next major version of Workload Characterization (WC) Tool (1.3) at all three labs; this version improved the demand/estimates interface and the reporting interface for additional computing resource usage and machine utilization reports
- Installed 1Gb networking infrastructure based on Interconnection Security Agreement (ISA) design and Memorandum of Understanding (MOU)
- Integrated the STAT lightweight debugger into all capacity platforms; supported it in Open MPI
- Developed a new, shared memory backing facility framework to be used by Open MPI's run-time system for memory footprint reduction

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- Retired the Purple compute system and its associated file system, including degaussing and disposing of over 17,000 classified disk drives, as well as dismantling and disposing of more than 1500 nodes and 200 racks; all trays and cable under the floor were also removed to clear space for Sequoia infrastructure
- Upgraded the facility control and monitoring system for the Terascale Simulation Facility (TSF) to improve resiliency and operation in preparation for Sequoia
- Deployed an unclassified general-purpose computing on graphics processing units (GPGPU) cluster for hybrid programming investigation and production simulation; deployed a classified visualization cluster (Muir) in preparation for Sequoia visualization requirements
- Engineered several security enhancements: 1) the HPC enclave is a network and system security implementation aimed at allowing sharing of expensive file system and archival storage resources among distinct user sets without compromising data security; 2) an additional one-time password service based on CRYPTOCARD technology was quickly delivered when the RSA (part of EMC corporation) announced a token security issue; this will be used to further strengthen Enclave security; and 3) off-site disaster recovery backups now use tape encryption for increased security of data at rest and in transit
- Hosted, on behalf of NNSA/ASC and the Office of Science (SC)'s Advanced Scientific Computing and Research (ASCR), *The 4th Workshop on HPC Best Practices: Power Management*, with participation from major HPC centers; the workshop report is available on the Web⁶

LANL Accomplishments for Facility Operations and User Support

- Under the New Mexico Alliance for Computing at Extreme Scale (ACES, with SNL), successfully completed Level 2 milestone 3952, *Cielo capability computing platform integration readiness* by delivering the Cielo platform and integrating it into the classified network, accrediting it for classified computing, completing Capability Computing Campaign 1 (CCC1); started CCC2 and upgraded to 1.37 petaFLOP/s; the computational environment usage model was developed, and it maps the Cielo capabilities to the tri-lab ASC Computing Environment (ACE) version 8.0 requirements
- Deployed Typhoon, a 106 teraFLOP/s cluster with 13,312 processors, on the secure network for capacity workloads
- Added capability to support tri-lab Cielo users, including a toll-free support line, online documentation on the classified network, integrated LANL/SNL HPC trouble ticket tracking system, and streamlining of the account granting process

⁶ <https://outreach.scidac.gov/pmbp>

- Installed an automated power and temperature monitoring system that has contributed to achieving a Power Usage Effectiveness (PUE) rating of 1.3 at the SCC
- Deployed Roadrunner with a proposal-based system similar to the CCC process for running the system at 60 percent for stockpile workloads and 40 percent for weapons science

SNL Accomplishments for Facility Operations and User Support

- Successfully completed the FY10 and FY11 Multi-Site Target (Milestone 3682 and Milestone 3941) to provide reliable, quality service, and access to any NNSA laboratory from any NNSA-designated computing user facility, independent of the location of the computing resource being utilized
- Supported National Security computing needs (Cielo del Sur, Cray XE6 platform) and added additional file system and tape storage capacity to the National Security Computer Center (NSCC)
- Retired the Red Storm computer system, transferred all data files to archive locations locally and remotely to LLNL and LANL, disposing of obsolete disk systems and redeploying many of the resources to support the NSCC mission
- Collaborated in solar energy work to gain experience for future computing facility design decisions
- Operated the TLCC1 capacity computing systems and supported the tri-lab TLCC2 selection and procurement process
- Deployed a new tiered user support model for SNL's ASC resources and integrated with LANL ASC user support to support the Cielo platform joint service desk

Academic Alliances

California Institute of Technology (Caltech) Accomplishments

- Completed, deployed, and analyzed a parallel implementation of the Parallel Optimal Transportation Meshfree method, the Center's Lagrangian solver. The performance of the implementation was analyzed with the aid of a benchmark test consisting of an 1100-0 aluminum cylinder striking a rigid wall at velocity of 300 m/sec.
- Continued development by the Computational Fluid Dynamics group of a Eulerian formulation for solid dynamics within the AMROC AMR framework. Several verification investigations were performed to provide confidence that the methodology produces sensible results.
- Completed design of the database schema for simulation archiving, and created a prototype Web-based explorer of the simulation archive. Major improvements in Mystic, the Center's massively parallel optimization framework, have eliminated dependence on third-party optimizers and enabled focus on development of optimization algorithms that are well-suited to the demands of Caltech's UQ program.

- Progressed significantly to describe higher Z atoms (p-block) in the electron force field, eFF, while retaining the cost-efficient spherical Gaussian wave packet representation of electrons.

Purdue University Accomplishments

- Conducted, under normal operation conditions, integrated static and dynamic pull-in simulation of the fluid-structure-electrostatic response of the Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM) microelectromechanical systems switch.
- Developed a first-of-its-kind, multiple-trap depth model for dielectric charging by integrating atomistic simulations of amorphous silicon nitride defect energy levels with a macroscale model for dielectric charging. The model was implemented in PRISM's MEMOSA software and UQ simulations were conducted to identify model sensitivities.
- Completed a full-system simulation of fluid-structure-electrostatics interaction, metal-dielectric contact, dielectric charging and creep using PRISM's coarse-grained switch model.
- Developed a mesoscale model for metal-dielectric contact based on atomic force microscopy (AFM) scans of the rough dielectric surface. The model is among the first to incorporate realistic descriptions of asperities and their role in rough-surface contact in microelectromechanical systems.
- Developed a creep model using long-term creep measurements performed at Purdue. The model was incorporated in MEMOSA and in the coarse-grained switch model. Bayesian calibration of creep model constants and UQ of the long-term (>500 hour) creep response were conducted, and may be the first of their kind to incorporate probabilistic models in creep prediction.

Stanford University Accomplishments

- Created a novel method to balance errors and uncertainties applied to a simplified full system, which showed importance of model-form uncertainties
- Implemented a viscous Reynolds-Averaged Navier Stokes (RANS)-solver in Domain-Specific Language (DSL) Liszt, which showed excellent parallel scaling up to 2048 cores
- Applied UQ (aleatory and epistemic) to full-system reacting flow simulations with modeled heat-release
- Demonstrated use of manufactured solution for code verification of high-speed reacting flow
- Showed that steady-state simulations can predict unstart, which was confirmed by in-house QMU experiment

University of Michigan Accomplishments

- Completed an assessment of predictive capability employing calibration

- Implemented integrated metrics for UQ with image data
- Published paper documenting capabilities and verification of the CRASH code
- Obtained near-optimal scaling of PDT radiation transport code to >16,000 cores
- Performed experiments to obtain early-time data needed for predictive science studies

University of Texas, at Austin, Accomplishments

- Completed first propagation of uncertainty in the full-system simulation using calibrated single physics
- Completed adjoint-enabled hypersonic flow code with coupled multi-physics
- Completed single physics calibrations across all modeling domains
- Added additional single physics and initial multi-physics manufactured solutions to University library for method of manufactured solutions
- Developed and implemented parallel multilevel modeling capabilities to facilitate Bayesian inversion on problems with multi-modal distributions and/or high dimensionality

IV. Product Descriptions by the National Work Breakdown Structure

The National Work Breakdown Structure (nWBS) divides ASC into five subprograms. IC, PEM, and Verification and Validation (V&V) are sub-programs presented in Volume 1 of the *ASC Implementation Plan*. CSSE and FOUS are presented, here, in Volume 2 of the *ASC Implementation Plan*.

WBS 1.5.4: Computational Systems and Software Environment

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of NNSA. This sub-program strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms, which include capability, capacity, and advanced systems. Along with these powerful systems that ASC will maintain and continue to field, the supporting software infrastructure that CSSE is responsible for deploying on these platforms includes many critical components, from system software and tools, to input/output (I/O), storage and networking, to post-processing visualization and data analysis tools, and to a common computing environment (CCE). Achieving this deployment objective requires sustained investment in applied R&D activities to create technologies that address ASC's unique mission-driven need for scalability, parallelism, performance, and reliability.

WBS 1.5.4.1: Capability Systems

This level 4 product provides capability production platforms and integrated planning for the overall system architecture commensurate with projected user workloads. The scope of this product includes strategic planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research for future systems.

WBS 1.5.4.1 Systems Requirements and Planning (LANL)

The Systems Requirements and Planning project covers all aspects of program and procurement planning for current and advanced systems and strategic planning for

supporting infrastructure. The major focus is to define requirements and potential system architectures for advanced systems platforms that meet ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts. In FY12, this project will focus on the project management of the Cielo system and additional systems. The focus in this project also includes the execution of DOE Order 413.3, using the NNSA Office of the Chief Information Officer Project Execution Model for information technology investments.

In FY11, the Cielo system was delivered to provide the tri-lab “user facility” capability machine. This included the completion of the ASC Level 2 milestone, *Cielo capability computing platform integration readiness*. This activity included the delivery and acceptance of the machine, completion of the milestone for system integration readiness for Cielo, and the delivery of capability compute cycles through the CCC1 and CCC2 periods.

Planned activities in FY12:

- Plan platforms as part of the exascale program
- Identify facility, power, memory, and file system requirements for exascale systems
- Participate in site-wide planning for power upgrades for future systems
- Plan infrastructure to support pre-exascale and exascale systems
- Participate in and develop the Science Partnership for Extreme-Scale Computing (SPEC) partnership between LANL, SNL, and ORNL, and associated activities

Expected deliverables in FY12:

- Updated platforms plan for LANL
- Requirements gathered for HPC
- Plans developed for future platforms, including dependencies on facilities or other areas
- Cielo critical decision (CD) 4 completed

Preliminary planned deliverables in FY13:

- Continue platform planning in conjunction with the other ASC areas

WBS 1.5.4.1 Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (SNL, LANL)

The Cielo capability computing platform is a project under the ACES. ACES is a joint collaboration between LANL and SNL defined under an MOU to provide a user facility for capability computing to the NNSA weapons programs in support of stockpile stewardship, to develop requirements and system architecture for ASC capability systems requirements definition, architecture design, procurement, key technology development, systems deployment, operations, and user support.

The architecture and design of Cielo is optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

Cielo is the primary platform that supports the ASC CCC. In FY10 it replaced the Purple platform at LLNL and is sited at LANL and operated by ACES. Cielo provides 1.37 peak petaFLOP/s with over 140,000 compute cores, and 10 petabytes of storage. Over 6,000 of the cores are dedicated to visualization and data services activities with connections to the SNL and LLNL sites.

In FY11, SNL and LANL deployed Cielo. The deployment consisted of two phases, the first being 1.03 peak petaFLOP/s in quarter 2 and then 1.37 peak petaFLOP/s in quarter 3. An early user period started in January 2011, which transitioned into the first campaign, CCC1, in February and continued through June 2011. CCC2 is expected to start in July and last through December 2011.

Planned activities in FY12:

- Complete the Cielo system integration
- Complete the ASC Level 2 milestone, *Cielo Production Capability Readiness*; complete the milestone review and write the report
- Transition Cielo platform to operations for full support of ASC CCCs

Expected deliverables in FY12:

- Cielo CCC2 and CCC3 campaign production capability cycles delivered to tri-lab users on Cielo
- Level 2 milestone, *Cielo Production Capability Readiness*, which indicates full production readiness of Cielo that is in use for CCCs

Preliminary planned activities in FY13:

- Support CCCs for tri-lab weapons programs

WBS 1.5.4.1 Alliance for Computing at Extreme Scale Architecture Office (LANL, SNL)

The primary objective for the ACES Architecture Office is to define requirements and potential system architectures for platforms that meet future ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts, such as a DOE exascale effort, and provides project management support for those efforts.

The ACES architecture office will coalesce mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design process. The ACES architecture office will also identify, in collaboration with the computer industry, critical technology gaps for future production capability systems, and support new technology development projects.

For FY12, in addition to the activities described above, two initiatives will form the focus of the Architecture Office, the Cray Interconnection Network D&E project and a new Advanced Memory Technology project with Micron Technology.

The Cray D&E is to ensure Cray's future interconnect, code named Pisces, addresses the needs the ASC applications in the pre-exascale timeframe. This will be accomplished by using co-design and simulation to influence the functionality and architecture of the Pisces interconnect.

The Micron cooperative agreement will examine the potential to extend Micron's Hybrid Memory Cube (HMC) memory technology foundation to dramatically lower the energy of data movement in exascale systems. It is well established that future platforms will have their power budgets dominated by data movement, and particularly DRAM access, which can run into hundreds of pico-Joules (pJ) per bit. Work in both architecture and technology is required to push the access energy to the sub-10 pJ/bit threshold required by exascale system requirements.

In FY11, a characterization of the Gemini high-speed network was used in Cray's XE6 architecture and hence Cielo. Cray delivered benchmark results and a performance report.

Planned activities in FY12:

- Research potential enhancements to architecture of the Pisces network interface chip to better support MPI
- Research possible enhancements, based on performance characterizations, to the Pisces router architecture
- Perform a comparative study between commodity InfiniBand interconnects and Pisces
- Work with Micron to establish the technology parameters for a 2018 memory system based on HMC technology
- Simulate Mantevo proxy applications using SST to evaluate the architecture of the logic layer of a memory cube and functionality that could be exported to the memory system
- Begin working with processor vendors to facilitate adoption of HMC technology

Expected deliverables in FY12:

- Reports on potential Cray Pisces high-speed interconnect (network interface chip (NIC) and router) enhancements to enable better support for ASC applications such as MPI
- Report on Pisces router enhancements based on simulations and performance characterizations
- An initial assessment of architectural alternatives to Micron HMC technology in support of mini-apps

Preliminary planned activities in FY13:

- Investigate possible enhancements, based on performance characterizations of real hardware, to the Pisces router architecture
- Analyze and prioritize recommendations for new capability and functionality to integrate into the Micron HMC architecture

WBS 1.5.4.2: Capacity Systems

This level 4 product provides capacity production platforms commensurate with projected user workloads. The scope of this product includes planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include the procurement and installation of capacity platforms.

WBS 1.5.4.2 Production Planning and Integration (LLNL)

The LLNL ASC strategy for capacity computing is to leverage industry advances and open source software standards to build, field, and integrate Linux clusters of various sizes into classified and unclassified production service. The programmatic objective is to dramatically reduce overall total cost of ownership of these “capacity” systems relative to best practices in Linux cluster deployments today. This objective strives to quickly make these systems robust, useful production clusters under the coming load of ASC scientific simulation capacity workloads.

A contract has been placed for a large amount of capacity computing resources to be delivered at all three sites over the FY11 and FY12 fiscal years. This tri-lab procurement will include options to procure a specific number of scalable units (SUs) each quarter from the third quarter of FY11 through the third quarter of FY12, for a total commitment of 54 SUs, with options for an additional 62 SUs over four quarters. The procurement will include Level 3 support and an onsite parts cache for each site.

This contract will provide a substantial increase in capacity compute resources for NNSA at all three labs. Clusters will have a common architecture and run the CCE software stack at each site. Options for clusters with GPGPU technology have been accommodated.

In FY11, LLNL led the tri-lab in the release, evaluation, and selection for the TLCC2 contract. LLNL also led the tri-lab planning and management for the TLCC2 project.

Planned activities in FY12:

- Receive, install, and integrate the TLCC2 SUs assigned to LLNL
- Ensure complete system delivery and ensure maintenance is included

Expected deliverables in FY12:

- TLCC2 SUs delivered to LLNL, operating in production computing environment, and running the TOSS stack

- TLCC1 systems at LLNL, operating in production computing environment, and running the TOSS stack

Preliminary planned activities in FY13:

- Provide production support for TLCC1 and TLCC2 systems
- Monitor computer industry developments for opportunities to enhance capacity computing and the associated infrastructure

WBS 1.5.4.2 Production Planning and Integration (LANL)

The Capacity System Planning and Integration project plans for new ASC capacity systems and integrates them into LANL classified, unclassified, and open computing networks. Included in this project is support for the ASC capacity system acquisition strategy and provision for requirements that help to achieve the strategy.

In FY12, the main focus will be on production support of capacity systems at LANL and planning for future TLCC architectures and integration, including the deployment of TLCC2. A major focus will also be on planning for TLCC2 and other capacity platforms and the infrastructure necessary for a successful deployment. Timely implementation and acquisition of site preparation, file systems, networks, archival storage, and applications development all need to be done in concert with capacity systems integration.

In FY11, LANL personnel participated in the evaluation and selection of TLCC2 capacity system vendor. LANL also continued evaluating market trends for capacity computing and understanding industry directions. Finally, Typhoon was added as a capacity system for classified computing

Planned activities in FY12:

- Provide production support of LANL capacity systems (Redtail, Hurricane, Typhoon, Lobo, Mapache, and Turing)
- Install TLCC2 capacity clusters
- Install and enable applications development on hybrid TLCC2 capacity clusters (GPGPU-enhanced TLCC2 clusters)
- Evaluate effectiveness of TLCC2 clusters for running large capacity jobs

Expected deliverables in FY12:

- Site preparation for TLCC2 SUs
- Initial production of TLCC2 capacity systems
- Installation and integration of capacity systems on the SU model to meet programmatic requirements for resources into FY12 and beyond

Preliminary planned activities in FY13:

- Provide production support for all capacity systems

WBS 1.5.4.2 ASC Capacity Systems (SNL)

The purpose of the ASC Capacity Systems project is to support the acquisition, delivery, and installation of new ASC capacity systems.

The project is supported by analysis of SNL's portfolio of application needs for capacity computing systems within the context of the higher integrated ASC platform strategy of capability, capacity, and advanced systems. Efforts include definition of requirements for TLCC systems and collaboration with the CCE product, with respect to a common software stack for new and existing capacity systems.

In FY11, SNL maintained and operated TLCC1 common capacity clusters with the CCE selected stack and tools in SNL's production computing environment. SNL supported the TLCC2 procurement process, including requirements gathering, architecture requirements, vendor selection, and preparation of a Statement of Work. SNL supported successful completion of the CSSE tri-lab Level 2 Milestone 3953, *TLCC2 Contract Awarded*. SNL supported the acquisition of the TLCC2 systems and provided testing on early delivery systems at LLNL.

Planned activities in FY12:

- Continue maintenance and operation of TLCC1 systems in SNL's production computing environment that are running the TOSS stack: Glory (2 SUs), Whitney (2 SUs), and Unity (2 SUs)
- Integrate CCE software stack and tools deployed on TLCC clusters
- Support tri-lab TLCC2 capacity systems delivery and installation
- Support deployment of the tri-lab TLCC2 systems with appropriate synthetic workload and acceptance testing; integration into each laboratory's environment; and transition from initial limited availability (LA) to GA
- Support the SNL FOUS Level 2 Milestone, *Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment*

Expected deliverables in FY12:

- TLCC1 systems operating in SNL's production computing environment running the TOSS stack
- Deployment of the tri-lab TLCC2 systems with appropriate synthetic workload and acceptance testing; integration into each laboratory's environment; and transition from initial LA to GA

Preliminary planned activities in FY13:

- Continue maintenance and operation of TLCC1 and TLCC2 systems in SNL's production computing environment that are running the TOSS stack

WBS 1.5.4.3: Advanced Systems

This level 4 product provides advanced architectures in response to programmatic, computing needs. The scope of this product includes strategic planning, research, development, procurement, testing, integration and deployment, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research, and the investigation of advanced architectural concepts and hardware (including node interconnects and machine area networks) via prototype development, deployment, and test bed activities. Also included in this product are cost-effective computers designed to achieve extreme speeds in addressing specific, stockpile-relevant issues through development of enhanced performance codes especially suited to run on the systems.

WBS 1.5.4.3 Petascale Application Enablement (LLNL)

The Petascale Application Enablement project enables advanced application work to develop benchmarks for new platforms, such as Sequoia, and to adapt current codes to the expected new architectures. A primary target of this project is investigating ways to improve application thread and I/O performance for future many-core platforms. The project team efforts include both direct application work and benchmark development and testing.

The project features detailed knowledge of performance analysis and evaluation. Core competencies include knowledge of computer architectures, system software, application coding strategies, and their interactions. The project has a proven record of optimization of key applications.

In FY11, LLNL evaluated expected performance for Sequoia benchmark applications and kernels from key science applications on the Sequoia prototype hardware and simulators. LLNL also continued investigations into threading performance issues with a particular focus on novel hardware features anticipated in Sequoia. The project also began investigations into the impact of node mapping strategies on Cielo and BlueGene systems. Performance improvements were realized on current architectures for key ASC applications, with a particular focus on IC.

Planned activities in FY12:

- Continue vendor interactions with respect to Sequoia application performance requirements
- Initiate science and weapons code testing on Sequoia and continue this activity on the Sequoia initial delivery (ID) system (Dawn) and Sequoia simulators, especially enhancing single-node thread performance
- Investigate opportunities for thread-parallel performance in production applications
- Identify mechanisms to enhance science and weapons code I/O performance

Expected deliverables in FY12:

- Evaluation of science applications and key benchmarks on Sequoia ID (Dawn)

Preliminary planned activities in FY13:

- Perform science and weapons code testing on the Sequoia system
- Continue with focus on code improvement opportunities

WBS 1.5.4.3 BlueGene/P and BlueGene/Q Research and Development (LLNL)

The BlueGene/P and BlueGene/Q R&D project is a multi-year R&D partnership between NNSA and DOE/SC with IBM on advanced systems. It targets the development and demonstration of hardware and software technologies for 1 petaFLOP/s and 10-petaFLOP/s systems. The BlueGene/P hardware is based on an extension of the highly successful BlueGene/L architecture with more cores per node, faster nodes, more memory, faster interconnects, and larger system scalability. The software approach to BlueGene/P is open-source collaborative development between IBM research, Linux Technology Center, the IBM Engineering and Technology Services Division, Argonne National Laboratory, and the ASC tri-labs. In FY08, a BlueGene/P system was delivered to Argonne. In FY09, a BlueGene/P system was delivered to LLNL as the first system (the Sequoia ID (Dawn)) in the Sequoia procurement. Follow-on BlueGene/Q system design targets a 20-petaFLOP/s system at the end of the contract.

This project incorporates requirements from the DOE laboratories, especially Argonne and the NNSA tri-labs, to have input into design choices and system testing for microprocessors, node architectures, and interconnects. The DOE laboratories also provide critical input on software, ensuring appropriate capability and features for the design target.

In FY11, LLNL led the Argonne/LLNL partnership in its interactions with IBM on the topic of BlueGene/Q. LLNL investigated compiler capabilities for utilization of Quad floating point (QPX), thread level speculation (TLS), and TM. Software and hardware simulation capabilities were used to determine Figure of Merit estimates for the Sequoia benchmark applications.

Planned activities in FY12:

- Continue technical interaction with IBM on hardware issues and software development
- Support Sequoia final phase 3 system build and early science period through technical interactions with IBM on hardware and software issues

Expected deliverables in FY12:

- Prototype of the new parallel runtime and the associated compiler changes of OpenMP (“cleansheet”); this work will involve fairly substantial changes to the compiler, including enhancements to the inter-procedural analysis phase; the ultimate

goal of this project is to reduce OpenMP overheads to such an extent that fine grain parallelism can be exploited effectively, allowing many more threads to be harnessed

- Threading performance enhancements
- Investigatory report on active storage alternatives for BlueGene/Q architecture
- Identification of SIMD architecture extensions (QPX) that would improve performance and/or broaden the range of applicability, and that are sufficiently general to merit consideration for incorporation into the architecture
- Ninety-six racks of Sequoia installed and integrated
- Early users on Sequoia hardware

Preliminary planned activities in FY13:

- Continue support of early science period
- Transition Sequoia to classified environment

WBS 1.5.4.3 Hyperion Test Bed (LLNL)

With the extreme demands for capacity computing, the I/O requirements of petascale applications for Sequoia, and the need for improved scientific data management capabilities, it is clearly apparent that emerging breakthrough technologies need to be tested in a large-scale environment such as Hyperion. The Hyperion Test Bed project will work with an expanded set of Hyperion vendor partners in the next phase of the Hyperion project to evaluate innovative node architectures, networks, and alternative storage solutions. Hyperion will continue to be a unique and critical resource for the functionality, performance, stability, and scalability testing of system software.

In FY11, LLNL added the following new capabilities: deployed new storage and networking hardware for the Hyperion Data Intensive Test Bed, including high performance flash memory and traditional disk storage systems, and with our partners, began the planning for technology refreshes in FY12.

Planned activities in FY12:

- Continue ongoing software and hardware functionality, performance, stability, and scalability testing and evaluation by the Hyperion community
- Evaluate data-intensive computing architectures for use in Hyperion phase 1 technology upgrade
- Evaluate the innovative storage and storage area network (SAN) architectures deployed in the Hyperion Data Intensive Test Bed

Expected deliverables in FY12:

- Technology upgrade and refresh of Hyperion phase 1 SU
- Deployment of the next generation Hyperion storage environment
- Cloud computing test bed into Hyperion environment

- Hyperion moves to the Livermore Valley Open Campus (LVOC) network and is available to a broader community of users

Preliminary planned activities in FY13:

- Employ technology upgrade and refresh of Hyperion phase 2 SUs
- Perform ongoing software and hardware scalability testing and evaluation by the Hyperion community

WBS 1.5.4.3 Advanced Systems Technology Research and Development (SNL)

The Advanced Systems Technology R&D project works to help overcome some of the bottlenecks that limit supercomputing scalability and performance through architecture and software research. Beginning in FY12, the focus will be on enabling technologies for co-design.

The project will address a critical need for experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy applications, enable application performance analysis with Mantevo proxy applications, support the Heterogeneous Computing and Programming Model R&D projects, and for SST validation efforts. These test bed systems will be made available for “test pilot” users that understand the experimental nature of these test beds. At the present time, it is more important to explore a diverse set of architectural alternatives than to push large scale. Discussions are already underway with Intel, AMD, IBM, Micron Technology, and other U.S. computer companies regarding ASC interest in obtaining early access to experimental architecture test beds. These partnerships will establish a strong foundation for co-design activities that can influence future hardware designs.

In FY11, the project focused on 1) advanced memory technologies, 2) advanced high-speed interconnects, and 3) runtime power management techniques. The high-speed interconnect research created a reference implementation for a next generation shared memory version of the Cray XT interconnect protocol. The runtime power research gained an understanding of the power requirements of ASC codes at large scale. These latter two activities have reached sufficient maturity that the efforts are being expanded into their own projects under the I/O and Networking and System Software and Tools products for FY12.

Planned activities in FY12:

- Procure two or more heterogeneous architecture test beds
- Procure advanced memory test beds

Expected deliverables in FY12:

- Mantevo proxy applications modified to match the capabilities of these new architectures

- Validation data for comparison with SST architectural simulations driven by these same Mantevo proxy applications
- Feedback to industry collaborators on the impact of specific hardware capabilities for future systems

Preliminary planned activities in FY13:

- Develop experimental runtime software to support mapping of application data structures to heterogeneous node architecture

WBS 1.5.4.3 Heterogeneous Computing (SNL)

The Heterogeneous Computing project will develop capabilities that facilitate ASC applications' ability to take advantage of heterogeneous architectures, including GPGPUs and other many-core processors. There are clear challenges for coupled, multi-physics-based simulation incorporating unstructured meshes and implicit solution methods. This project will work in tight alignment with co-design efforts and mini-app development to explore performance on available heterogeneous architecture test beds. Programming models and associated runtime support for portable hybrid parallelism and data locality/placement must be developed to efficiently exploit the diverse set of many-core processors proposed for these architectures. In addition, fine-grained, architecture-aware load balancing of work must be explored and utilized.

These multi-physics coupled applications have a higher level of heterogeneous parallelism that is not currently being exploited. This project will explore SNL's ability to support parallel invocation of coupled applications and efficient management of the shared data. Both physics and engineering applications critical to the nuclear weapons enterprise must be able to take advantage of all levels of parallelism to benefit in the context of exascale computing.

This project will build upon researchers' expertise to improve application performance and portability to the next generation of architectures for scaling applications to a billion-way parallelism. This project has close ties to the heterogeneous architectures test beds in SNL's Advanced System Technology R&D project.

This project is new for FY12.

Planned activities in FY12:

- Research programming models and associated runtime support for portable computational kernels and their data structures on heterogeneous architectures, including GPGPUs and other many-core processors
- Partner with the IC sub-program element to support development of computational kernels that incorporate heterogeneous-parallel capability
- Extend existing finite element Mantevo mini-apps to demonstrate programming models and optimize performance of computational kernels on heterogeneous architectures

- Research fine-grained load balancing techniques at the node level with support scalability for applications
- Research programming models and frameworks to exploit heterogeneous and hierarchical parallelism within applications

Expected deliverables in FY12:

- Mini-app prototype demonstrating significant performance improvement for the Mantevo miniFE application through the use of improved data structures and kernels on heterogeneous architecture test beds

Preliminary planned activities in FY13:

- Partner with mini-app developers and multi-physics applications to incorporate developed programming models to fully and optimally exploit parallelism in heterogeneous architectures

WBS 1.5.4.4: System Software and Tools

This level 4 product provides the system software infrastructure, including the supporting operating system (OS) environments and the integrated tools, to enable the development, optimization, and efficient execution of application codes. The scope of this product includes planning, research, development, integration and initial deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include system-level software addressing optimal delivery of system resources to end-users, such as schedulers, custom device drivers, resource allocation, optimized kernels, system management tools, compilers, debuggers, performance tuning tools, run-time libraries, math libraries, component frameworks, other emerging programming paradigms of importance to scientific code development and application performance analysis.

WBS 1.5.4.4 System Software Environment for Scalable Systems (LLNL)

The System Software Environment for Scalable Systems project provides system software components for all the major platforms at LLNL, research and planning for new systems and future environments, and collaborations with external sources such as the platform partners, especially IBM and Linux vendors. This project covers system software components needed to augment Linux and required proprietary OSs that function in a manageable, secure, and scalable fashion needed for LLNL ASC platforms.

This project includes work on developing, modifying, and packaging the TOSS, and developing scalable system management tools to support the OS and interconnect (for example, TOSS and InfiniBand monitoring tools), as well as the resource management environment (Moab and Simple Linux Utility for Resource Management (SLURM)) to queue and schedule code runs across LLNL systems. LLNL uses TOSS on all of its Linux clusters. This project also funds approximately 60 percent of the manpower

required to develop, deploy, and maintain TOSS. The funding LLNL receives for its portion of CCE TOSS funding accounts for 40 percent of the effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in section 1.5.5.6 of this document.

In FY11, LLNL added the following new capabilities: deployed TOSS 1.4 (based on RHEL 5.5), released SLURM 2.2, developed TOSS 2.0 (based on RHEL 6) including integration support for TLCC2 hardware, and modified SLURM (v2.3) to support Sequoia system.

Planned activities in FY12:

- Continue ongoing development/support for TOSS software
- Deploy TOSS 2.0 (based on RHEL 6)
- Deploy InfiniBand SAN support for Sequoia
- Continue ongoing development and support of Moab and SLURM
- Port SLURM to Sequoia system

Expected deliverables in FY12:

- TOSS 2.0 deployed
- InfiniBand SAN support for Sequoia deployed
- SLURM 2.3 released to accompany the TOSS 2.0 release
- SLURM port to Sequoia in production

Preliminary planned activities in FY13:

- Continue ongoing TOSS software development and support
- Investigate alternate network topologies for clusters and SANs
- Continue ongoing development and support of Moab and SLURM
- Participate in Exascale System Software Working Group

WBS 1.5.4.4 Applications Development Environment and Performance Team (LLNL)

The Applications Development Environment and Performance Team (ADEPT) project provides the code development environment for all major LLNL platforms, supports user productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The elements of the development environment covered by this project include, but are not limited to, compilers, debuggers, performance and memory tools, interfaces to the parallel environment, and associated run time library work.

In FY11, the ADEPT project included a high priority focused activity within ADEPT, entitled Scalable Applications Preparation (SAP), that will ensure a wide range of ASC applications run effectively on Sequoia, including exploiting its novel hardware and software. It also includes ADEPT responsibility for the Level 2 milestone 3854, *Scalable applications and preparation outreach for Sequoia*. In addition, ADEPT has deployed production-quality versions of portions of the scalable code correctness suite, with a particular focus on the STAT. Further, ADEPT has begun planning for exascale systems, including an evaluation of the most critical emerging issues and strategies to address them.

Planned activities in FY12:

- Provide, maintain, and refine major platform code development environments
- Coordinate integrated design code scaling for Sequoia as part of SAP activities
- Explore novel threading strategies and performance as part of SAP activities
- Continue design and development of exascale development environment approaches
- Explore new programming models for future capability systems
- Interact with the ASC code teams and vendors to improve software products

Expected deliverables in FY12:

- Deployment and testing of the Sequoia development environment and characterization of the performance on the system
- Sequoia system applications development tools environment improvements
- Refinement of development environment plan for exascale system
- Deployment of a full production version of highly scalable code correctness tool suite
- Certification of the TLCC development tools environment for new TOSS releases

Preliminary planned activities in FY13:

- Continue code development environment support on all LLNL ASC platforms
- Refine code development environment for existing and future systems
- Continue exploration of programming models for exascale systems

WBS 1.5.4.4 Co-Design through Mini-Applications (LANL)

The Co-Design through Mini-Apps project will build a co-design process through the collaborative creation of patterns, strategies, and abstractions for the implementation and optimization of scientific applications and algorithms on emerging hardware architectures. The primary vehicle for this process will be a suite of open source proxy applications (mini-apps), derived from and feeding back into ASC integrated code teams. ASC code teams will provide the requirements for the study based upon application domains of interest to ASC. The specification documents and reference implementations produced will act as the basis for most of the work in this project, ensuring that it is

targeted directly to ASC code needs. In the long run, the multi-year results of this project will be used as input to new ASC exascale code development efforts.

The project will track the development of next-generation hardware architectures and study both computational and data movement patterns represented by the chosen mini-apps developed in conjunction with ASC IC code developers. We will also investigate novel hardware data transformation techniques prototyped using reconfigurable hardware.

Work done directly with integrated IC teams will continue to enhance existing performance and tools, and to evolve current codes towards next-generation architectures by investigating improved mesh data structures, new strategies for AMR, performance improvement, data locality, compressed data structures for materials. In addition, we will explore increased software abstraction, emerging programming models, and DSLs.

This is a new project for FY12. Work will exploit the Darwin R&D test bed (a GPU-based system with approximately 100 nodes, each with 48 cores and NVIDIA or AMD accelerators) that was delivered at the end of FY10 and deployed in FY11.

Planned activities in FY12:

- Work with ASC code teams to select primary areas for mini-app development, and assist code teams in developing specification documents and reference codes
- Study mini-apps, their effective implementation on next-generation architectures, and work with ASC code teams to identify and implement useful abstractions for multi-physics applications on extreme-scale computing systems
- Continue work on re-factorization of ASC codes, moving new code capabilities into production
- Continue to build and leverage collaborations with SNL and LLNL to address scientific computing at extreme scale through synergistic mini-app efforts

Expected deliverables in FY12:

- Mini-apps characterizing ASC IC physics and usage
- A prototype data transformation engine using reconfigurable hardware
- Improved performance metric and tracking capabilities delivered to code projects

Preliminary planned activities in FY13:

- Continue tracking the evolution of next-generation architectures and the associated impact on exemplar applications
- Explore more mini apps and more architectures; extend lessons learned from analysis of mini-apps to realistic applications
- Extend data transformation engine with user programmable functionality via a domain specific data translation language

WBS 1.5.4.4 Programming Models for the Next-Generation Scientific-Computing Environment (LANL)

The goal of this project is to study emerging hardware trends and their impact on programming abstractions/models, the overall software development tool chain, and run-time systems for scientific-computing environments. LANL's goal is to develop a set of technologies that will assist in the development of the next-generation of application codes as well as extend the lifetime of current codes.

Motivated by the ongoing revolution in computer architectures, this project will experiment with existing techniques and explore new models. It will also support compile- and run-time infrastructures needed to provide the scientific-computing community with the necessary tools to address the challenges of developing software on extreme-scale, highly concurrent systems. Work will be guided by leveraging both existing and evaluation hardware combined with the needs of the IC and PEM workloads, and common computational and communication patterns.

In FY11, the project developed technologies for enhanced program analysis and transformation and developed new approaches to understanding and managing the complex memory hierarchy of emerging multi-core processors.

Planned activities in FY12:

- Collaborate with the Co-Design through Mini-Apps and Visualization and Data Analysis projects to investigate the impact of using existing application programming interfaces (APIs) to program next-generation architectures for both physics and data-intensive applications
- Develop techniques and tools for the analysis and potential transformation of code structures given the dynamic behavior of select applications (the goal is to assist in the identification of target code elements for acceleration on new architectures, and explore how these findings can be leveraged to define new programming models)
- Develop techniques and tools for the analysis of data movement through the memory hierarchy and its impact on application performance and scaling (including the I/O subsystems); explore how these findings can be leveraged to define new programming models
- Participate in and support the activities of the ASC CSSE exascale planning working group on programming models

Expected deliverables in FY12:

- Identification of potential limitations of current parallel programming models/APIs
- Example implementations of programming abstractions for emerging processor and system architectures
- Implementation of static-, compile- and run-time analysis tools that can be leveraged to improve developer productivity on both current and next-generation architectures

Preliminary planned activities in FY13:

- Advance software development technologies, including improved programming models (in terms of both productivity and performance) and improved developer feedback for effective/efficient software development on future platforms
- Deliver software developed in FY12 to meet the FY13 Level 2 milestone, *Programming Models and Data Analysis Environments for Extreme-Scale Systems*, that will provide a programming model and environment that addresses the challenges of programming advanced architectures

WBS 1.5.4.4 High Performance Computing Systems Research (LANL)

HPC systems research is a broad-reaching project focusing on near to long-term research of all the components needed to support a rich environment for very large-scale applications. Systems research bridges the gap between hardware and programming model, and requires tight collaboration in supporting the development of programming models, tools, visualization/analytics, and system software aspect of I/O.

The project is currently focused on resilient system services, soft-error resilience, system support for DISC, and performance modeling/simulation.

Resilient system services focus on developing a vehicle to investigate resilient, dynamic, distributed, scalable services for large-scale systems and providing an interface to programming models so that ASC applications can access these features on current and future hardware.

Investigations of soft-error resilience will focus on creating an accelerated testing environment for soft error profiling using a virtual machine (VM) approach to inject faults while actively running real ASC codes.

The performance simulation investigation is currently focused on application I/O patterns used on parallel file systems under realistic application loads.

In FY11, the project investigated a dynamically resilient approach to system services and established new efforts in resilience and fault tolerance. Basic Distributed File System Simulation capability (FileSim) was added using the scalable SimCore modeling framework that was also made interoperable with SNL's SST.

Planned activities in FY12:

- Use VMs to investigate application susceptibility to soft errors
- Investigate core partitioning on high core count nodes
- Investigate resilient protocols for supporting dynamic, distributed, and scalable services
- Increase fidelity of FileSim through adding protocol details of the object storage device of the Panasas file system and incorporate a communication network simulator into FileSim

- Perform extensive validation tests against large-scale I/O traces obtained from Cielo
- Participate in and support the activities of the ASC CSSE exascale planning working group on system software

Expected deliverables in FY12:

- Prototype of soft error susceptibility profiling (that is, gprof for soft errors)
- Dynamic distributed library for building fault tolerant services
- Validation report of FileSim tests on Cielo and user documentation of FileSim
- Fault Tolerance at Extreme Scale (FTXS) Workshop organization

Preliminary planned activities in FY13:

- Extend resilient system services to other types of networks and services; investigate application to future many-core systems
- Extend the toolset for soft error susceptibility
- Increase level of detail model for file system simulation to near-cycle accurate level and expand model to other file systems; use FileSim in procurement process

WBS 1.5.4.4 Application Readiness (LANL)

The Application Readiness project addresses issues with an application's production-run readiness on current and incoming computing systems at LANL. Working with subsystem teams such as systems management, file systems and I/O, archive, and tools, the Application Readiness team identifies causes of unexpected behavior and deploys fixes in production. The project goal is that system users are able to make productive use of the systems with their applications to solve their problems.

The project provides production problem solving (create small problem reproducers, identify cause, consult with the relevant technical experts to find a solution, and verify the deployed solution), periodic stress testing/regression of production machines, new software version regression testing, system configuration verification and software stack deployment with real user applications and metrics, and analysis/profiling.

In FY11, LANL expanded the skill set of the team with an additional hire and by including applications engineers from vendors to support the Cielo computing platform.

Planned activities in FY12:

- Stand up applications on in-coming production systems, notably deployment of a new TLCC2 system
- Support users of Cielo, Mapache, and other recently integrated systems by tackling hard-to-diagnose problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Work with other projects including Co-Design through Mini-Apps, Programming Models, and Code Strategies for Emerging Platforms to assist code development

teams with migration towards the use of hybrid programming to exploit the cell processors on Roadrunner and on the GPU-based elements of a new TLCC2 system

- Provide support to production applications on capacity platforms
- Assist system management personnel with problem investigation and resolution

Expected deliverables in FY12:

- Science/code team application runs and regression testing completed to standup the TLCC2 system and its move toward an operational state
- GPU-enabled codes running on TLCC2

Preliminary planned activities in FY13:

- Continue to assist users with migration towards the use of hybrid programming
- Assist system management personnel with problem investigation and resolution

WBS 1.5.4.4 Software Support (LANL)

The Software Support project works to establish a strong development and analysis tool capability for current and next generation HPC platforms, including parallel capabilities. It is focused on working with the HPC tool community and vendors to identify, plan, and integrate tools into production environments and establish a solid support structure. The project supports the incremental improvement of tools driven by ASC strategic needs, including cross-laboratory partnerships and external collaborations that focus on performance tools required for programming model support.

Capabilities include development of a strategic plan for tools based on current and next generation platform planning, integration with HPC community tool development efforts, and tool development and production integration capability.

In FY11, LANL leveraged the work of the CCE project to further establish a tri-lab collaboration work space, supported the scalable debugger work with TotalView Technologies, and began instituting tool workshops between users and tool providers to provide user training and develop requirements.

Planned activities in FY12:

- Continue building an MPI support capability by engaging the community support model; focus will be on Open MPI for development needs
- Continue performance analysis support capability; focus will be on O|SS integration in application analysis and new analysis needs
- Increase memory analysis and MPI performance/correctness tool support
- Continue building a stronger debugger support capability; work with TotalView on tool deployments and enhancements; work with LLNL debugger capability through STAT and subset debugging

- Continue involvement with DOE NNSA and DOE/SC laboratories in exascale tool planning and development programs
- Develop the Component-Based Tool Framework (CBTF), a common, scalable framework to ease implementation, integration, and deployment of software development tools

Expected deliverables in FY12:

- MPI analysis tools for memory, performance, and correctness deployed
- O|SS deployed on production platforms, including Cielo and TLCC2
- Interaction with community and vendors to increase MPI scale capability
- Initial prototype of CBTF
- Software support and development capability strategy plan

Preliminary planned activities in FY13:

- Continue building an MPI support capability by engaging the community support model
- Continue building a stronger debugger and performance analysis support capability

WBS 1.5.4.4 Software and Tools for Scalability and Performance (SNL)

The Software and Tools for Scalability and Performance project supports system software R&D to address scalability of future computational systems. An important aspect is providing lightweight services and functionality that do not compromise scalability and therefore performance. The focus starting in FY12 will be on emerging critical areas for HPC systems. Power has become a first-order design constraint for future supercomputers. SNL will expand upon work in data collection techniques that provided new insight into understanding power requirements of ASC applications. The increasing complexities of node memory layouts will be addressed by developing system-level memory locality correction techniques using a lightweight kernel (LWK). LWK technology will also be used to explore a revolutionary mechanism for providing shared libraries in a highly scalable manner.

This project has provided LWK OSs and highly scalable run time software for ASC's capability-class systems. This year, a quantitative analysis of the results of that research will be a key component of a Level 2 milestone, *Demonstration of a Legacy Application's Path to Exascale*. The team has an understanding of dynamic shared libraries and the issues associated with deploying them at high scale on capability systems. This project has close ties to the heterogeneous architectures test beds under Advanced System Technology R&D activity in WBS 1.5.4.3.

In FY11, Red Storm system software R&D and experience was applied to the Cielo deployment. A file system configuration was designed for dynamic shared library support

on Cielo. The results of the Cielo 6X acceptance tests were analyzed to develop an analytical method for predicting performance of future platforms.

Planned activities in FY12:

- Implement power-aware infrastructure into an HPC compute node OS
- Investigate frequency modification techniques at scale to achieve maximum application power efficiency
- Investigate future CPU advanced power management architecture features
- Identify application, OS, and hardware power feedback loop requirements
- Design and implement system software-level mechanism for continuously optimizing placement on systems with non-uniform memory access (NUMA), based on observed application behavior during runtime
- Specify requirements for hardware and software interfaces to measure component power on commodity boards

Expected deliverables in FY12:

- Lightweight compute node OS capable of controlling frequency based on individual applications' optimal value
- Analysis of performance improvement due to system software-level optimization of data placement on a Cielo platform
- Lightweight VM monitor supporting multi-core guest OSs embeddable in LWK host OS
- Prototype implementation of peer-to-peer caching scheme for dynamic shared library support

Preliminary planned activities in FY13:

- Continue efforts to raise system-software awareness of application behavior and adapt accordingly
- Draft API to provide feedback loop of power usage, needs, and available resources

WBS 1.5.4.4 Resilience (SNL)

Resilience has been identified as a critical issue for a viable exascale computing system. Resilience in this context refers to the ability of applications to run correctly to completion even in the face of failures in hardware and system software. Aggregate component failure rates are expected to outpace disk I/O rates at exascale, such that the overhead of checkpoint/restart will outweigh its value as a failure mitigation strategy.

The Resilience project's goal is to develop scalable techniques, mechanisms, and enhancements to hardware and system software to enable increased understanding of system behaviors and application-system interactions and ultimately to reduce the risk of failure at the application level. SNL will use this information to build fault/degradation

models that include root cause and fault propagation understanding in the context of current large-scale ASC HPC platforms. These models will be used for developing failure detection/prediction mechanisms for invoking mitigating response mechanisms and influencing future architectures. This work will enhance the HPC communities' ability to characterize system state, understand system faults, and develop mitigation mechanisms. Four key core capabilities associated with this effort are: 1) scalable, real-time failure and system characterization, 2) failure and system modeling, 3) failure detection and prediction, and 4) development of effective response mechanisms.

In FY11, SNL developed and applied information collection, analysis, and visualization services to Cielo platform to extract conditions of interest from system logs and hardware monitors. In particular, these were development and deployment of Splunk modules on RedSky and Cielo, a new log file analysis and visualization capability, and Web-based mechanisms for interaction with distributed storage and analysis system.

Planned activities in FY12:

- Collaborate with Cray on XE6 platform data collection; utilize enhanced data collection capabilities to support high-fidelity run-time evaluation of system performance and to develop fault models to be used in enhancing resilience and guiding new platform development

Expected deliverables in FY12:

- Deployment of scalable data collection infrastructure on Cielo or similar platform for characterization and modeling of observed system faults and performance, and development of relevant fault models
- Investigation of system-wide complications on analysis and response such as intra-application resource contention and cascading effects of response

Preliminary planned activities in FY13:

- Deploy production hardened resilience tools in large-scale production environments to utilize run-time data collection in conjunction with relevant fault models to initiate mitigating response mechanisms on a large-scale platforms

WBS 1.5.4.4 System Simulation and Computer Science (SNL)

The SST is a suite of tools that enables multi-scale computer architecture simulation to support HPC software/hardware co-design. The SST consists of a core set of components that enable parallel discrete-event simulation; high-fidelity networking, memory, and processor components; and coarse-grained simulation components that capture essential elements of machine performance with low computational cost. Future HPC systems and the applications designed to utilize them are impacted by a variety of considerations, including scalability of applications, ease-of-programming, memory and network latencies becoming more imbalanced relative to computation rates, data corruption and its propagation, frequency of interrupts, power consumption, and overall machine cost. SST is designed to allow each of these parameters to be explored, permitting the consideration of a broad space of potential architectural and application/algorithmic designs. The goal

is for the SST components to be extended and enhanced by a community of simulator developers. An even larger community is expected to be the users of SST, including algorithm developers, architecture designers, and procurement team members.

In FY11, support for MPI in SST/macro and DUMPI was greatly expanded. An effort to apply UQ to discrete event simulation in the context of SST was initiated. The quality and performance of coarse-grained network congestion models was improved. Features were added to DUMPI to facilitate the transfer of trace files from classified computing resources to external partners. The M5 processor simulator was integrated to support higher fidelity modeling of X86 and ARM processors. Several system technology models were integrated to analyze system reliability and thermals, and energy usage models were integrated into SST/micro. More flexible memory models were integrated to analyze advanced capabilities. Software Quality Assurance (SQA) practices were greatly improved with automated regression testing and generation of distribution and documentation artifacts.

Planned activities in FY12:

- Generalize DUMPI to support I/O and non-MPI communications traces
- Implement fault injection and power models in SST/macro and improve network congestion models in SST/macro
- Interface with ASC exascale co-design centers to define a co-design process
- Compile a database of SST simulator results for growing portfolio of Mantevo proxy applications, and perform validation experiments with advanced architecture test beds
- Develop and analyze alternative performance, energy, thermal, and reliability management policies
- Implement UQ algorithms in a discrete event simulation framework

Expected deliverables in FY12:

- Public SST code release incorporating new capabilities and analysis models
- Library for temperature/energy aware partitioning
- Tool for hardware/software co-simulation

Preliminary planned activities in FY13:

- Improve coarse-grained processor models
- Integrate simulation and analysis efforts for stochastic variability
- Develop high-level cost models to add to technology library
- Develop objection function for co-design optimization

WBS 1.5.4.4 Application Performance Analysis with Mantevo Proxy Applications (SNL)

The purpose of the Application Performance Analysis with Mantevo Proxy Applications project is to continue the development of the Mantevo mini-apps and to facilitate their use as a tool to judge the performance of current and next-generation technologies and platforms.

The current capabilities of the Mantevo Project can be found on the Web.⁷ In FY12, SNL will develop new benchmarks and metrics that are relevant to how ASC applications perform. A related concern is the development of objective functions to drive the optimization process that is an element of the co-design methodology. A key goal is the development of new objective functions for co-design that are closely tied to SNL's new metrics for exascale computing.

In FY11, a new mini-app called miniGhost was developed. It was used as part of the Cielo Phase 2 acceptance testing, with execution across core counts numbering over 100,000. At this scale, miniGhost stressed the architecture in ways that are expected to be of special interest at exascale, including collective communication and irregular mesh computations. In addition, miniXyce was deployed and miniFE was used by industry and academia for evaluation of new processor technologies, including Intel's MIC family and nVidia's Fermi GPGPU.

Planned activities in FY12:

- Use Mantevo mini-apps with SST and in experimental architecture test beds to understand the trade-offs among some often competing priorities, including time to solution, energy requirements, memory utilization, I/O requirements, and resilience performance
- Survey current methods for measuring capability, performance, energy efficiency, and productivity
- Investigate and characterize metrics that are of importance to ASC applications

Expected deliverables in FY12:

- Documentation for the original developers of Mantevo mini-apps (primarily application developers supported by IC) and of the changes SNL or industry collaborators are making to provide better performance on advanced architectures
- Definition of co-design objective functions

Preliminary planned activities in FY13:

- Define initial metrics, methods, and benchmarks for measuring pre-exascale systems
- Perform initial evaluation of current platforms using mini-apps

⁷ <https://software.sandia.gov/mantevo/>

WBS 1.5.4.4 Scalable, Fault-Resilient Programming Models (SNL)

ASC has identified several key software development challenges on the road to exascale. While the several-orders-of-magnitude increase in parallelism is the most commonly cited of those, hurdles including drastically shortened mean times to interrupt, increased imbalance between computational capacity and I/O capabilities, silent errors, and complex hardware architectures. Various programming and execution models have been proposed to address these issues, but before ASC applications can use these, detailed insight will be required into the level of porting effort, scalability, and fault tolerance characteristics of these approaches.

Current efforts on architecture simulation, performance prediction and analysis, and mini-app analogues of full-scale ASC codes will be used to identify required algorithmic changes and to quantify the performance and fault resilience characteristics of proposed exascale programming models. Efforts will focus on both on-node execution models required to support upcoming heterogeneous architectures in an “MPI+X” parallel environment, and on alternate software development approaches with potential to support greater levels of parallelism and fault resilience than the cooperative message passing SPMD approach commonly associated with MPI.

This effort is new in FY12 but will leverage several existing efforts funded by ASC and ASCR that are addressing heterogeneous/GPGPU platforms, mini-apps and skeleton applications for ASC codes, advanced data types for fault resilient computing, and advanced OS functionality for exascale platforms. Additionally, this effort will be integrated with current simulation efforts to provide performance prediction and analysis of alternate programming models. This project has close ties to the heterogeneous architectures test beds under Advanced System Technology R&D activity in WBS 1.5.4.3.

Planned activities in FY12:

- Integrate with current efforts on alternative programming models, including data structures for fault-oblivious data-centric computing, evaluation of probabilistic computing methods, and advanced OS functionality for efficient data sharing across coherency domains
- Evaluate and update Mantevo mini-apps to support explorative programming model research; select and develop metrics for performance and development costs appropriate for ASC application development
- Integrate and extend the current simulation capabilities to support mapping applications to suitable programming models; predict performance and fault-resilience of refactored applications at the exascale

Expected deliverables in FY12:

- Mantevo mini-app converted to directed acyclic graph scheduling
- Trace file collection run on current HPC platforms and performance analysis using SST/macro

Preliminary planned activities in FY13:

- Optimize performance of modified Mantevo mini-app using SST
- Validate simulator performance predictions using Cielo and other current ASC platforms
- Create initial metrics for development effort and performance profile analysis for non-MPI programming models

WBS 1.5.4.5: Input/Output, Storage Systems, and Networking

This level 4 product provides I/O (data transfer) storage infrastructure in balance with all platforms and consistent with integrated system architecture plans. The procurement of all supporting subsystems, data transfer, storage systems, and infrastructures occurs through this product. The scope of this product includes planning, research, development, procurement, hardware maintenance, integration and deployment, continuing product support, quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include high-performance parallel file systems, hierarchical storage management systems, storage-area-networks, network-attached storage (NAS), and high-performance storage system (HPSS) or future hierarchical storage management system disks, tape, robotics, servers, and media. This product also includes relevant prototype deployment and test bed activities. Projects and technologies in the advanced networking and interconnect areas include networking and interconnect architectures, emerging networking hardware technologies and communication protocols, network performance/security monitoring/analysis tools, and high performance encryption and security technologies.

WBS 1.5.4.5 Archive Storage (LLNL)

The Archival Storage project provides long-term, high-performance, archival storage services to ASC customers. This includes a collaborative software development effort (currently, HPSS) between the tri-labs, Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, and IBM, as well as deployment and support of archival storage software and interfaces for tri-lab ASC customers on unclassified and classified networks. It includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and media, ongoing technology refresh, and data stewardship. HPSS provides scalable, parallel, archival storage interfaces and services to the tri-labs.⁸

A world-class array of hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms, thereby increasing computation. This includes disk arrays, robotic tape subsystems, servers, SANs, networks, and petabytes of tape media, enabling high-speed parallel transfers into a virtually unlimited data store.

⁸ See <http://www.hpss-collaboration.org/index.shtml>.

In FY11, LLNL completed the ASC Level 2 milestone 3855, *Deploy high-performance storage system quota system*. HPSS R7.4 and Redundant Array of Independent Tapes (RAIT), version 0, were developed, HPSS R8.1 design was finalized and development began. HPSS R7.3.2 was deployed featuring LTO-5 support. LLNL selected, procured, and deployed upgrades of LTO tape drives and media on new Linux clusters, and of Oracle Titanium tape drives and media to existing Linux clusters. Obsolete hardware (including all AIX mover machines) was decommissioned.

Planned activities in FY12:

- Finalize development and testing of HPSS R7.4, a minor release with conversion to 64-bit architecture, support for IPv6, and repack existing small files into aggregates
- Finalize development and testing of RAIT.0, a limited minor release with RAIT, providing higher reliability cost-effectively
- Develop HPSS R8.1, featuring an architecture utilizing distributed, parallel components and partitioned metadata to better position HPSS for exascale
- Support the ASC Level 2 milestone, *HPC Enclave*, for instantiating security zones, by procuring archive solutions or enhancing security measures for new zones
- Deploy a later release of HPSS (R7.4.x or R7.5.x) in production archive environments
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support, and maintenance of all archival storage hardware and media, customer and interface support, ongoing tech refresh, and data stewardship

Expected deliverables in FY12:

- Release of HPSS R7.4 and R7.5, and RAIT.0 to limited user base

Preliminary planned activities in FY13:

- Develop and release HPSS R8.1
- Begin planning for production deployment of HPSS R8.1

WBS 1.5.4.5 Parallel and Network File Systems (LLNL)

The Parallel and Network File Systems (NFS) project provides for the development, testing (feature, capability, performance, and acceptance), procurement, integration, and ongoing support of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included are the continuing development and support of Lustre as a fully featured file system for the range of ASC capability and capacity platforms, the deployment and support of global parallel file system (GPFS), the deployment and support of ubiquitous NAS services for home, project, and scratch space, and the I/O support of various programming interfaces for parallel I/O.

This project deploys and supports Lustre and GPFS file systems for ASC platforms as well as high-availability NAS file systems for home and project space, and scratch space for serial capacity clusters. It actively works with the OpenSFS Lustre development

community to add Lustre file system scalability and reliability enhancements required by ASC platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

In FY11, LLNL ported the ZFS file system to Linux and fielded a proof-of-concept ZFS-based Lustre file system in preparation for the Sequoia platform. Lustre deployment teams procured, deployed, and operated a new multi-petabyte file system in classified and unclassified environments.

Planned activities in FY12:

- Deploy ZFS-based Lustre Sequoia file system
- Provide Lustre file system performance and scalability enhancements in support of Lustre 2.1 release
- Work closely with application teams and Sequoia developers on implementation and performance of the Lustre client on Sequoia hardware
- Expand production Lustre file systems in preparation computational increases
- Change NAS architecture for LLNL HPC enclave
- Investigate NAS acceleration and caching strategy in support of large-scale job startup
- Maintain and support GPFS, NAS, and Lustre parallel file systems, including middleware and higher level I/O libraries for users

Expected deliverables in FY12:

- Deployment of 50-petabyte Sequoia ZFS-based Lustre file system
- Development and deployment of Lustre 2.1
- Replacement of existing architecture with HPC NAS enclave solution

Preliminary planned activities in FY13:

- Complete development of porting layer with Lustre, allowing multiple file system back-ends
- Deploy initial pNFS NAS file systems

WBS 1.5.4.5 Networking and Test Beds (LLNL)

The Networking and Test Beds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect subsystems in support of current and future systems and environments. This work relies heavily on an adequately provisioned test bed, skilled staff, and collaborations with vendors.

This project will test various hardware and software components to quantify the features, performance, reliability, security, and interoperability of the products and broader technology base. The information acquired as a result of this project will be used to help determine an integrated architecture and resultant procurements for these subsystems.

In FY11, host adaptor problems continued to plague performance as larger file systems taxed the I/O infrastructure. The test bed and related tools isolated and reproduced the problem such that the vendor could update their firmware. The worldwide economic downturn delayed the introduction of much of the technology LLNL expected to test. LLNL was able to evaluate and test Quad Data Rate (QDR) InfiniBand switches and cables. LLNL prepared the request for quote for Sequoia SAN procurement.

Planned activities in FY12:

- Perform research and testing for technologies and products for interconnects, local area networks (LANs), WANs, and National Security Agency (NSA) Type 1 encryptors, file system servers, clients and disks, with special focus on emerging DataCenter Ethernet switches, additional features in InfiniBand, other interconnect and 10GigE-related technologies, and 10GigE NSA Type 1 encryptors in support of future ASC petaFLOP/s systems
- Continue study of developing load balancing and multi-path routing in support of congestion avoidance for previously listed networks
- Apply testing results to optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems

Expected deliverables in FY12:

- Initial installation of Sequoia SAN network

Preliminary planned activities in FY13:

- Continue to leverage tri-lab activities in I/O-related hardware and software, and seek to improve the reliability, performance, and manageability of the I/O subsystems in production
- Research and test to determine which technologies and products should be considered for insertion into production to meet the growing I/O performance and capacity requirements
- Continue network improvements and track emerging network technology, especially related to congestion

WBS 1.5.4.5 Archival and File Systems (LANL)

The Archival and File Systems project provides design and development of archival services, and high-performance file systems and I/O infrastructure for the ASC program. Capabilities supported include online file systems such as the NFS complex and enterprise-wide supercomputer file systems, GPFS development, deployment and management, scalable I/O (SIO) middleware development and support, interconnect technology development and deployment, SAN development and deployment, and archive.

The file systems element of the project provides end-to-end, high-performance networking and SIO infrastructure for the ASC program. Successfully meeting the ASC programmatic milestones requires carefully balanced environments in which the I/O

infrastructure scales proportionally with increased ASC platform capabilities and application data needs. As the program moves toward exascale areas, these efforts will improve the scaling or programmability of the I/O in ASC applications for current and future large-scale machines. Current areas of investigation are Parallel Logged File System (PLFS), scalable indexing, burst buffer architectures, and scalable metadata.

The archival element of the project develops the HPSS (HPSS is a collaborative software development effort between the tri-labs, Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, and IBM) and parallel storage interface (PSI) software (LANL) for the purpose of supporting ASC customers from LANL, LLNL, and SNL, including all aspects of HPSS and PSI planning, development, deployment, and support. HPSS is software that manages petabytes of data on disk and robotic tape libraries. HPSS provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape. HPSS uses cluster, LAN, and/or SAN technology to aggregate the capacity and performance of many computers, disks, and tape drives into a single virtual file system of exceptional size and versatility. This approach enables HPSS to meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored.

In FY11, LANL added the following new capabilities: PSI 2011 release, redefined release strategy for HPSS, refined requirement documents for HPSS 8.0 release, and 7.3 release of HPSS.

Planned activities in FY12:

- Develop PLFS
- Explore scalable NFS
- Verify client/server versions of Panasas and assist with file system stand-up of incoming machines
- Support the Applications Readiness project with any I/O issues and general production needs
- Finalize requirements for design, develop, and test 2012 release of PSI
- Develop requirements for 2013 release of PSI
- Continue support for the production archive

Expected deliverables in FY12:

- FY12 Level 2 milestone, *Application Deployment of a Quick Parallel Log-Structured File System Capability*
- Tested deployment of scalable NFS service to replace current large appliance approach
- PLFS tested and deployed with LISTIO and index reduction threads, and prototyped with a scalable indexing service
- PSI 2012 release; releases of HPSS 7.4 and HPSS 7.5

Preliminary planned activities in FY13:

- Release HPSS 8.x beta and HPSS 8.1
- Continue planning, development, deployment, and support for the production archive

WBS 1.5.4.5 Archival Storage (SNL)

The Archival Storage project represents SNL's participation in the DOE HPSS Consortium development project. HPSS provides the archival storage solution for ASC systems and is in direct alignment with ACES.

SNL's role in the HPSS project is to collaborate with tri-lab developers to design, implement, and test solutions that meet ASC requirements for all three labs.

In FY11, SNL completed and released HPSS version 7.3. This release added high availability and LT05/T10KC tape support along with creating an mkhpss command line. RAIT was also developed for HPSS.

Planned activities in FY12:

- Provide support for PFTP, HACL, and HPSSD
- Test and release HPSS 7.4, which will include support for 64-bit codes, IPv6 transport protocol, and significant improvements to the repack subsystem and dynamic reconfiguration
- Design, test, and release HPSS 7.5, which will include full hpssadm functionality for command line, alternate copy, real-time monitoring support for PVL/PVR services
- Design HPSS 8.1, which will include multi-core servers and metadata; identify impacts and new requirements based on the new version 8 architecture
- Extend capabilities of HPSS version 8 by adding support for content validation, user-defined attributes, high-availability heartbeat and failover, and multilevel security

Expected deliverables in FY12:

- HPSS 7.4 release
- HPSS 7.5 design, test, and release

Preliminary planned activities in FY13:

- Provide support for PFTP, HACL, and HPSSD
- Release HPSS 7.6
- Beta test HPSS 8.1

WBS 1.5.4.5 Scalable Input/Output Research (SNL)

The SIO project provides support for I/O library and file systems on existing petascale ASC platforms as well as critical R&D to provide I/O capabilities on future exascale platforms. The research performed in this project directly addresses two vital concerns

for I/O on exascale platforms: scalable parallel file systems and technologies for integration of computation and analysis.

To improve file-system scalability and resilience, the SIO project is developing a file system that decentralizes management of devices to support a high degree of heterogeneity within a system of inherently unreliable networks and storage devices. The central components of this peer-to-peer-like system are “smart” servers that have access to a variety of different local and remote media (for example, disk, NVRAM, memory, and tape) and are pervasive throughout the computing platform. These servers directly handle I/O requests, initiate third party transfers, or replicate the data as needed.

The SIO project is also developing technologies for integration of computation and analysis. This software will have tremendous impact on I/O for petascale and future systems because it allows for the creation of integrated scientific workflows that process data before the data reaches persistent storage, avoiding storage of transient data products. Current use of this technology includes data staging/caching to manage bursty I/O operations (for example, for checkpoints) and in-transit fragment detection for CTH, a co-design activity with the Scalable Data Analysis project.

In FY11, the SIO project completed work on the GPGPU RAID and deployed prototypes on a storage cluster at University of Alabama at Birmingham and at SNL for testing within the production environment. The GPGPU RAID work was also selected as a candidate for an R&D 100 award. Finally, the SIO group collaborated with the Trilinos project to form the Trilinos I/O Support (Trios) capability area. Trios provides a vehicle for deployment and testing of software products developed by the SIO research project.

Planned activities in FY12:

- Support I/O and file system issues for Cielo, TLCC, and test bed platforms, and address file system issues
- Continue activities for pNFS and POSIX I/O API extensions
- Perform R&D for the name server, storage server, and security modules for the peer-to-peer file system
- Perform R&D to support integrated computation and analysis as a co-design effort with Scalable Data Analysis (WBS 1.5.4.6)
- Perform R&D for simulation of disks, RAID, and NVRAM devices as a co-design effort with the system simulation project (WBS 1.5.4.4)

Expected deliverables in FY12:

- Prototype implementations of network, storage, and security modules for peer-to-peer file systems
- In-transit fragment detection service for CTH for the FY13 Level 2 milestone, *Data Co-Processing for Extreme Scale Analysis*
- SST simulation components for RAID disks and generic I/O servers (based on Lustre)

Preliminary planned activities in FY13:

- Perform large-scale tests of peer-to-peer file prototypes on Cielo
- Perform R&D for data-management services to support data provenance and efficient data movement for codes with integrated analysis

WBS 1.5.4.5 Scalable Interconnects for Extreme-Scale Tightly Coupled Systems (SNL)

The Scalable Interconnects for Extreme-Scale Tightly-Coupled Systems project will develop capabilities to enable performance and scalability of ASC applications on current and future high-performance interconnection networks on extreme-scale platforms. This project will concentrate on characterizing application requirements with respect to functionality and performance for intra-application data movement as well as application network transfers to external I/O services. This project will also provide a low-level network programming interface appropriate for current-generation network hardware as well as more advanced next-generation hardware with more sophisticated network interface capabilities and functionality. As applications explore alternative programming models beyond the current distributed memory MPI model, the low-level network programming interface must evolve to include the ability to provide very lightweight one-sided data transfer operations, while continuing to enable efficient two-sided message-based transfers. It is likely that this project will expand to include an analysis of network topologies, network interface hardware design and evaluation, optimized network transfer protocols, and system software support for advanced network interface operations.

This project will build on existing efforts surrounding the development of the next-generation Portals network programming interface and measurements of application sensitivity to network performance. It has close ties to the ACES project and the Cray network hardware design and engineering activity in WBS 1.5.4.1.

This project is new for FY12.

Planned activities in FY12:

- Continue to develop the Portals 4.0 network programming interface specification
- Measure application sensitivity to network injection bandwidth and network link bandwidth on a capability-class system
- Develop and support a reference implementation of the Portals 4.0 network programming interface
- Continue evaluation of Portals triggered operations to increase the performance of MPI point-to-point and collective communication operations
- Continue implementation and evaluation of Portals 4.0 for the Partitioned Global Address Space programming model
- Continue to participate in MPI forum activities and ensure that Portals supports the new capabilities being standardized

Expected deliverables in FY12:

- Release of an initial reference implementation of Portals 4.0 for InfiniBand networks
- An initial analysis of the sensitivity of applications to network injection bandwidth
- A study of Portals 4.0 triggered operations for MPI non-blocking collective operations

Preliminary planned activities in FY13:

- Publish a detailed analysis of the sensitivity of applications to network injection and link bandwidth for several important applications on a large-scale, capability-class system

WBS 1.5.4.6: Post-Processing Environments

This level 4 product provides integrated post-processing environments to support end-user visualization, data analysis, and data management. The scope of this product includes planning, research, development, integration and deployment, continuing customer/product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include tools for metadata and scientific data management, as well as general-purpose and application-specific visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data—the use of open-source foundations will continue to be an important strategy for development of shareable advanced techniques. The product must develop solutions to address interactivity, scaling, tri-lab access for petascale platforms, and data analysis techniques needed to support effective V&V and comparative analysis. Solutions for emerging platform architectures may in turn require customization and/or re-architecting of software to leverage hardware features. A continuing emphasis will be placed on tools for improving end-user productivity. The product also provides and supports infrastructure including office and collaborative space visualization displays, mechanisms for image data delivery, and graphics rendering hardware.

WBS 1.5.4.6 Scientific Visualization (LLNL)

The Scientific Visualization project conducts research and develops and supports tools for managing, visualizing, analyzing, and presenting scientific data. Research topics include topological analysis, particle visualization, and data compression techniques. Operational support for data analysis covers support of post-processing resources, including visualization servers, displays, and facilities. The visualization hardware architecture team engages in planning, test bed prototyping, testing of systems and components, and procurement and integration of new systems. Display efforts include support of high-resolution, high-performance display devices for theaters and collaborative use areas. The project installs, maintains, and consults on software visualization tools, and supports demonstrations on the PowerWalls. The project maintains unclassified and classified video production labs and consults on software such

as resource management tools, movie players, animation, and visualization packages. The project exploits the latest capabilities of clustering hardware, GPUs, and parallel storage systems. Hardware capabilities include three production visualization servers and several PowerWall clusters. A video display infrastructure drives PowerWalls and smaller displays. Visualization researchers continued to perform work in areas of topology, compression, and advanced data analysis techniques.

In FY11, the LLNL visualization team deployed a new architecture for driving PowerWalls with a smaller number of more powerful nodes. The team also deployed technologies for remote display, procuring and installing PC-over-IP technology. The team delivered a suite of visualization software on the new open computing facility (OCF) cluster, Edge, a 216-node cluster with GPUs.

Planned activities in FY12:

- Provide operational support for all visualization facilities, including supporting projection equipment and facilitating the use of the data analysis clusters and associated storage
- Support large-scale data analysis activities through software consulting and research activities
- Perform R&D in topological analysis, data compression, and particle visualization, and continue to mentor students in these areas
- Maintain and enhance the existing suite of tools and libraries developed by this project in support of PowerWall and visualization cluster usage
- Support ASC scientists through visualization and video efforts, including the support of PowerWall presentations and the creation of visuals and movies to support the presentation of scientific data

Expected deliverables in FY12:

- Smaller, more powerful PowerWall clusters deployed in ASC theaters
- Suite of analysis and visualization software deployed on visualization productions clusters, including GPU-related tools

Preliminary planned activities in FY13:

- Continue to enhance and maintain the data analysis hardware and software environment
- Leverage research efforts in visualization and data analysis and perform new research in data compression and topological analysis for scientific computing

WBS 1.5.4.6 Scientific Workflow and Data Management (LLNL)

In the age of petascale computing environments, the complexity and scale of the data management challenge is also reaching new heights. Commensurate with this is the difficulty of managing the simulation workflow, from problem definition to post-processing the data. The Scientific Workflow and Data Management project provides

users with powerful and time-conserving ways to access, search, compare, and archive large-scale scientific data, and new high-level tools for managing the simulation workflow. This is achieved through the development of production-quality applications that enhance data management capabilities and the creation of innovative interfaces to job monitoring and vertical application frameworks.

Hopper and Chopper are the principal products of this effort. Hopper is a Java-based file management tool that allows users to transfer and manipulate files and directories by means of a graphical user interface. Users can connect to and manage local and remote resources using all file transfer protocols supported by ASC computing centers. Chopper is the command line version of the tool, useful in particular for automated, “background” file manipulations initiated from within applications.

In the simulation workflow area, the existing Lorenz Web-based HPC application suite forms a foundation for providing new ASC-specific capabilities. Lorenz uses advanced Web technologies to make HPC more accessible, saving the user time while also helping the resources to be used more effectively.

In FY11, LLNL released versions 2.6 and 2.7 of Hopper and Chopper, featuring expanded support for data aggregation operations, a job resubmission capability, simple embedded content viewing for new file formats, a history editor, and numerous scalability and usability enhancements aimed at next-generation datasets.

Planned activities in FY12:

- Maintain and enhance the existing suite of tools and libraries developed by the scientific data management project

Expected deliverables in FY12:

- Version 3.0 of Hopper and Chopper, featuring a data transfer wizard that helps the user organize and optimize file transfers for peak performance, extended usability improvements, and additional cross-laboratory transfer capabilities
- A production release of ASC-relevant job monitoring and job interaction functionality through the Web-based Lorenz job management application
- A beta release of at least one vertical application of interest to ASC, using the Web-based Lorenz application portal framework

Preliminary planned activities in FY13:

- Continue to maintain and enhance the Scientific Data Management suite of tools and libraries based on user feedback and requirements imposed by a petascale computing environment
- Investigate the sharing of code components between Hopper and Lorenz

WBS 1.5.4.6 Production Visualization (LANL)

The purpose of the Production Visualization project is to provide LANL weapons designers with visualization systems research and support, and to provide analytic expertise to help LANL weapons designers utilize the full power of the hardware and

software infrastructure for visualization and data analysis developed and deployed by ASC, thus improving the physics understanding of their weapons simulations.

Technical staff members funded by Production Visualization assist in the design theater and the co-laboratories. The project also deploys within the design community in X Division a small group of individuals with expert knowledge in both visualization and weapons science to work directly with the designers.

The project's capabilities include the design and deployment of new visualization systems, briefing support, and support of large facilities, such as the CAVE and the PowerWall. The project provides assistance, training, and developing of new tools to work with these facilities. Development, deployment, and maintenance of any needed visualization corridor software is also provided by this project.

The project is responsible for the EnSight visualization and data analysis software, including maintaining the EnSight software installation laboratory-wide, providing local user support in the use of the software, and acting as a bridge between the LANL design community and the EnSight developers at Computational Engineering International.

In FY11, LANL added the following new capabilities: Viewmaster 2 deployment into production; Cielo visualization to LANL, SNL, and LLNL, Blue Room visualization facility; and visualization infrastructure at distances not before reached from the SCC to the NSSB.

Planned activities in FY12:

- Design and deploy new PowerWall
- Continue to support and maintain production visualization systems
- Provide contract management and requirements specification, including facilities, visualization cluster, and EnSight contracts
- Support and maintain the EnSight software and help direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International
- Work directly with designers in physics-based, iterative discovery process using the petascale visualization and data analysis enabled tool (EnSight) to promote new discoveries in weapons science by in programs such as the Thermonuclear Burn Initiative (TBI)

Expected deliverables in FY12:

- Upgraded PowerWall facility in production
- Visualization in production in the NSSB
- EnSight, PoP, and GMV upgrades
- Scientific visualization support for new users; new weapons science visualizations
- Other visualizations and data analysis products to address a variety of DSW-related activities, including the closing of outstanding SFIs

Preliminary planned activities in FY13:

- Continue to support and maintain production visualization systems
- Continue to promote new discoveries in weapons science by advanced applications of visualization and data analysis
- Continue to document the results with classified papers and publications on weapons science topics jointly co-authored with X Division designers

WBS 1.5.4.6 Visualization and Data Analysis Research and Development Project (LANL)

The Visualization and Data Analysis R&D Project develops new visualization algorithms and systems to meet capability requirements for ASC simulations. This work is required to address ASC workloads: massive data sizes, complex results, and the use of unique supercomputing architectures.

ASC simulations are currently producing massive amounts of data that threaten to outstrip the ability to visualize and analyze it. Therefore, it is important to understand how to triage data within the simulation as it is generated using techniques such as *in-situ* analysis for data reduction and visualization kernels that run on the supercomputing platform, including data analysis, visualization, and rendering methods.

Another important analysis mode involves storing reduced data for later interactive analysis and visualization (that is, post-processing analysis). Extreme-scale ASC databases, petabytes or larger, already exist and are growing in number and size. LANL's long-term objective is to build a post-processing analysis system that, within reasonable limits, can manage performance optimization tasks automatically, be fault-tolerant, provide a set of high-level serial and parallel programming primitives for carrying out complex queries and computations on the stored data, and provide fast enough execution to enable interactive, iterative discovery. Current approaches range from the database-driven (SQL), to parallel computation-driven (Hadoop/MapReduce), neither of which alone can meet the needs of the target community. LANL will build on these technologies, integrating the best of both approaches.

In FY11, LANL supported the exascale planning effort, leading a tri-lab visualization and analysis team to produce a white paper and presentation for the ASC community.

Software for multi-resolution visualization, parallel ray tracing on multi-core machines, feature extraction algorithms, and a parallel netCDF reader have been made available to the ASC community in the latest releases of Visualization Toolkit (VTK) and ParaView. Research was published, including work on parallel rendering on multi-core machines, science-based feature extraction, a data-intensive storage system that supports visualization, and an *in situ* statistical sampling data triage approach.

Planned activities in FY12:

- Explore the end-to-end integration of *in-situ* analysis and feature-extraction visualization to improve the understanding of massive scale results

- Develop a data-intensive post-processing visualization and analysis infrastructure that meets ASC programmatic needs by leveraging existing and new custom data intensive software
- Participate in and support the activities of the ASC CSSE exascale planning working group on visualization and data analysis

Expected deliverables in FY12:

- A portable, open-source supercomputing visualization, analysis, and rendering library based on the Thrust data parallel programming library
- Integration of statistical sampling and visualization/analysis to provide a mathematically rigorous approach to effectively reduce massive data sizes
- An *in-situ* data triage and reduction framework for analysis in coordination with SNL
- A report on how to optimize file reads/writes for efficient access during post-processing visualization and analysis

Preliminary planned activities in FY13:

- Deliver software developed in FY12 to meet the Level 2 milestone, *Programming Models and Data Analysis Environments for Extreme-Scale Systems*; this milestone will provide a data analysis environment that addresses the challenges of massive data, including *in situ* analysis and advanced data reduction techniques; these products will be a step towards enabling productivity and performance at exascale

WBS 1.5.4.6 Scalable Data Analysis (SNL)

The Scalable Data Analysis project provides production data analysis tools, R&D, and support for a range of ASC customers—from analysts and code developers to algorithm designers and hardware architects. In addition, this project provides deployment and support services that enable ASC customers to carry out data analysis on ASC capability and capacity systems, with a primary focus on delivering and supporting analysis capability for Cielo and subsequent ACES platforms. Current capabilities include scalable data analysis software released open source through ParaView and the VTK, *in-situ* data analysis library for coupling directly with running codes, and ensemble analysis R&D.

Partnering with ASC customers and other product areas, this project will build on the successful ParaView, co-processing library, and VTK products, and will continue R&D that advances these capabilities as needed for evolving next-generation architectures, ensuring that ASC's investment in data analysis and visualization will provide advanced capabilities on platforms from Cielo through future exascale architectures.

In FY11, both application and foundation technologies scaled to handle Cielo-sized dataset, so that ParaView was used to interactively visualize the largest AMR dataset to date. This is a successful extension of capability already in production use and demonstrates the robustness of the core technologies developed thus far for ASC. In

addition, an open source and *in-situ* analysis library based on VTK was released; it was incorporated into a range of codes from a variety of science domains.

Planned activities in FY12:

- Provide targeted analysis capabilities to promote an integrated user experience to improve productivity and insight into results, in conjunction with SNL codes
- Continue research into analysis of complex data ensembles, with an emphasis on visualizing and understanding uncertainty, both in inputs (for example, constituent models) and the resulting ensembles of output results
- Partner with customers to provide advanced analysis techniques through easily accessible, browser-enabled interfaces
- Provide ongoing technical and user support, including direct analysis support for capability-class systems
- Partner with I/O project to provide technologies for integrated computation and analysis

Expected deliverables in FY12:

- Advanced analysis capability in support of GA of the Cielo platform, including high-performance distance visualization and platform-specific optimization
- Continued development of ParaView, including domain-specific capabilities and scaling needed for analysis of Cielo-sized data and beyond
- Web-based analysis capabilities, including ensemble analysis tools incorporating uncertainty analysis capabilities
- Advanced analysis and visualization library for direct coupling with codes
- Ongoing releases and installations of ParaView

Preliminary planned activities in FY13:

- Analyze, *in-situ*, for the Level 2 milestone, *Data Co-Processing for Extreme Scale Analysis*

WBS 1.5.4.7: Common Computing Environment (Moved)

This level 4 product was moved at the beginning of FY12 to 1.5.5.6: Common Computing Environment.

WBS 1.5.5: Facility Operations and User Support

This sub-program provides both necessary physical facility and operational support for reliable production computing and storage environments as well as a suite of user services for effective use of ASC tri-lab computing resources. The scope of the facility operations includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities, and collaborations. FOUS also covers physical space, power and other utility infrastructure, and LAN/wide area network (WAN) networking for local and remote access, as well as requisite system administration, cyber-security, and operations services for ongoing support and addressing system problems. Industrial and academic collaborations are an important part of this sub-program.

WBS 1.5.5.1: Facilities, Operations, and Communications (Retired)

This level 4 product was retired at the end of FY10 and replaced with two new products, which began in FY11:

- 1.5.5.4: System and Environment Administration and Operations
- 1.5.5.5: Facilities, Network, and Power

WBS 1.5.5.2: User Support Services

This level 4 product provides users with a suite of services enabling effective use of ASC tri-lab computing resources. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. Projects and technologies include computer center hotline and help-desk services, account management, Web-based system documentation, system status information tools, user training, trouble-ticketing systems, and application analyst support.

WBS 1.5.5.2 Hotlines and System Support (LLNL)

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC computing resources for the tri-lab as well as academic and industrial collaborations. This project includes computer center hotline and help desk services, account management, Web-based system documentation, system status information tools, user training, incident management systems, and application analyst support. Services are provided to both LLNL users as well as users from external sites, including LANL, SNL, and the ASC Alliance sites.

This project provides accounts administration, technical consulting, and documentation and training to facilitate the effective use of LLNL HPC systems. An accounts specialist team provides all account management services necessary for users to obtain accounts and access LLNL HPC systems. This includes account creation and removal, bank allocations, token management and visitor tracking for foreign national users. The technical consultant team provides technical support to LLNL users to enable their effective use of LLNL HPC systems. Consulting services vary from helping new users configure their environment, assisting experienced users with optimization of codes, and supporting other Livermore Computing (LC) staff with monitoring of file systems, batch queues, and user environments. Extensive Web documentation, user manuals, technical bulletins, and training are provided to users via email, Web, and in-person training.

In FY11, LLNL implemented the Front Range problem management module in hotline operations. Phone messages now inform users that issues are being resolved and all open incidents associated with a problem are resolved when the problem is resolved.

Planned activities in FY12:

- Continue to implement information technology infrastructure library (ITIL) best practices (change management or configuration management)
- Consolidate and update HPC Web infrastructure and documentation
- Improve systems status tools for more effective user communication
- Provide ongoing support services for hotline operations, documentation, and training
- Evaluate hotline metrics on a regular basis to identify potential areas that would benefit from improved documentation and training

Expected deliverables in FY12:

- Sequoia and TLCC2 system documentation and training materials

Preliminary planned activities in FY13:

- Continue to implement ITIL best practices
- Implement Configuration Management Front Range module
- Provide ongoing support services, documentation, and training for hotline operations

WBS 1.5.5.2 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)

The Integrated Computing Network Consulting, Training, Documentation, and External Computing Support project is responsible for direct customer service for local and remote users of ASC/LANL resources, the development and delivery of documentation and training materials for ASC/LANL resources, usage statistics, and an administrative interface for the ASC tri-lab, Alliance users, and other external ASC/HPC users. The primary capabilities consist of user support services, operational metrics for an HPC environment on, for example, usage and availability, Web-page development to present

this information to system personnel and users, and the development of user documentation and training.

In FY11, LANL, in cooperation with SNL, added capability to support tri-lab users of Cielo, including a toll-free support line, online documentation on the classified network, integrated LANL/SNL HPC trouble ticket tracking system, and streamlining of the account granting process by eliminating LANL-specific training requirements.

Planned activities in FY12:

- Perform ongoing user support for users of ASC/LANL/ACES computing resources
- Identify major sources of end-user system outages using 2011 baseline data
- Develop comprehensive system software testing regime to proactively (before customers) detect system software faults and unintended consequences, with special emphasis on the Cielo environment
- Develop modern Distributed Data Services (DDS) middleware layer that builds on CCE projects for delivering local HPC system data

Expected deliverables in FY12:

- Data-supported analysis of key factors limiting end user compute availability
- Framework for continuous system software testing in place and functioning
- DDS middleware layer in place and processing at least three disparate data sources across capability and capacity platforms

Preliminary planned activities in FY13:

- Deploy an integrated program of system correctness testing to include GPU correctness
- Develop and initiate action plan to increase overall availability of compute cycles to end users
- Analyze system software fault rates via deployed testing framework from fiscal year 2012 data
- Develop new customer interfaces to deliver comprehensive job information using DDS middleware

WBS 1.5.5.2: User Support (SNL)

The User Support project provides user support and associated resources for SNL computing systems and tri-lab resources. User support activities focus on improving the productivity of the entire user community, local and remote, in utilizing the ASC HPC resources.

This project deploys and maintains the following SNL capabilities for user support:

- Coordination between user support activities and leadership in adopting ITIL principles and practices

- ITIL-compliant incident, problem, and knowledge management tool set
- Training facilities and equipment
- A Web portal for HPC-related information, real-time data, and documentation

In addition, this project provides the following user support capabilities in conjunction with other projects:

- A tiered user support structure (HPC service desk) that responds to SNL and tri-lab user requests received via phone, email, Web-based requests, and in-person visits
- The Synchronized Account Request Automated Process (SARAPE) tri-lab account provisioning Web-based tool
- Web-based, classroom, and one-on-one training
- Direct support in utilizing ASC resources

This project also funds the SNL user support team's involvement in collaborative efforts such as the Predictive Science Academic Alliance Program (PSAAP) and ACES.

In FY11, SNL added the following new capabilities: a production tiered user support structure (the HPC OneStop Service Desk); use of an ITIL-compliant Service Desk tool as a pilot for wider SNL corporate use; an upgrade to a new version of the incident management tool in conjunction with SNL corporate transition to the tool; processes and capabilities to maintain status and event information in the SNL HPC Web Portal.⁹

Planned activities in FY12:

- Provide coordinated, tiered user support for SNL's ASC resources and, in partnership with LANL, for ACES resources
- Support SNL capabilities such as issue tracking tools, training facilities, and the SNL HPC Web portal
- Support SNL applications for large runs on Cielo and other ASC platforms

Expected deliverables in FY12:

- Deployment of additional ITIL processes within the ITIL-compliant tool suite, including problem management
- Coordinated user training across SNL HPC OneStop user support sub-teams
- Improved user support tools, including enhanced collaboration and communications
- Reliable and responsive service to users in the ASC tri-lab computing environments

Preliminary planned activities in FY13:

- Continue to leverage ITL as a framework for improving the HPC OneStop service desk processes and practices

⁹ <https://computing.sandia.gov>

- Continue to enhance user support tools, capabilities, and infrastructure available to users and the user support team
- Provide assistance for use of emerging ASC platforms/architectures

WBS 1.5.5.3: Collaborations

This level 4 product provides programmatic support for collaboration with external agencies on specific HPC projects. This product also includes collaborations with internal or external groups that enable the program to improve its planning and execution of its mission.

WBS 1.5.5.3 Program Support (LLNL)

The Program Support project provides service to the ASC program. Program Support services include procurement and contracting, project management, and meeting support. These services are in support of both tri-lab and LLNL-only activities, including collaborations with academic, industrial, and other government agencies.

In FY11, the most significant change in the Program Support project was the completion of the TLCC2 procurement selection phase.

Planned activities in FY12:

- Continue management of the Sequoia contract and its associated R&D and D&E contracts with IBM
- Initiate, as needed, next-generation and architecture-related contracting activities
- Manage existing tri-lab contracts, including TLCC
- Support the annual Supercomputing conference, Predictive Science Panel (PSP) meetings, and other meetings and workshops such as the continuation of Best Practices Workshop series
- Support PSAAP collaborations
- Provide support to the ASC Federal program management office
- Participate in planning with the HPC community for a possible future exascale program in partnership with the DOE/SC; this may include activities such as workshops

Expected deliverables in FY12:

- Execution, if approved, of initial phases of next-generation procurement activities, as well as continuation of CCE and TLCC-related contract management

Preliminary planned activities in FY13:

- Continue FY12 contract management activities

WBS 1.5.5.3 Program Support (LANL)

Through the Program Support project, LANL provides support to the national program, both by providing resources and expertise to the Federal program office and by participating in coordination and integration activities for the tri-lab program.

In FY11, LANL hosted the PSP, which reformed under a new charter with a new panel in FY11. The new PSP charter focuses on providing technical feedback on the combined efforts of ASC and Science Campaigns for predictive science. In FY11, LANL also provided consultant support to the Federal program efforts to foster collaborations while building support within the predictive science community, LANL supported the PSAAP, provided support for the Supercomputing conference, and brought editorial and publishing support on board at LANL for supporting tri-lab publications and the ASC newsletter.

Planned activities in FY12:

- Alternate with LLNL in hosting the PSP; results will be incorporated into program plans and initiatives
- Provide consultant support to the Federal program management efforts to foster collaborations and build support within the predictive science community
- Support the PSAAP

Expected deliverables in FY12:

- PSP meeting
- ASC Principle Investigator Meeting
- ASC newsletter

Preliminary planned activities in FY13:

- Host the PSP
- ASC Newsletter

WBS 1.5.5.3 Program Support (SNL)

The Program Support project funds critical coordination and integration activities essential to the success of ASC. It is divided into two distinct parts: 1) provide ASC programmatic planning, reviews, and communications; and 2) SNL outreach to the other institutions and programs.

This capability is critical to the ASC SNL program integration, communication, and management within the laboratory and with the external community. A significant management and integration function in this project is captured in the SAIC contract that provides support for NNSA Headquarter (HQ) and SNL in communications and logistics. External advisory boards supported through this project also provide feedback to the ASC leadership team regarding the maturation of the predictive engineering sciences capability and the quality of SNL's computational science R&D. Support of external collaborations,

including the University Alliance program and the exascale initiative (with DOE/SC), is also included in this project.

In FY11, SNL added the following new capabilities: organized and hosted third Predictive Engineering Science Panel meeting in September 2011; supported two QASPR external review panel meetings; supported the Engineering Sciences External Advisory Board and the Computational Sciences External Advisory Board. These boards look at the quality of the ASC technical work as it underpins the development of predictive engineering sciences capabilities. SNL also supported current PSAAP collaborations, worked with HQ and tri-lab community to define future PSAAP program call and engagement with potential bidders, provided support to the ASC Federal program office with emphasis on the exascale initiative, developed and supported SNL ASC Web page and connection to NNSA page, organized a successful Nuclear Weapons Engineering Analysis Conference with participation from across the complex to discuss the most challenging engineering analysis challenges and potential solutions, supported the Supercomputing Conference and PSP meetings, supported the ASC executive committee; supported quarterly meetings of the ASC executive committee, and managed the new SAIC contract to provide various administration support to HQ.

Planned activities in FY12:

- Support external advisory panels for developing predictive engineering capability
- Lead NNSA booth at SC11 Supercomputing Conference
- Develop and begin implementation of advanced architecture plans

Expected deliverables in FY12:

- Support for HQ and the tri-lab to articulate the drivers and challenges for exascale computing
- External advisory boards for SNL's developing predictive engineering capability
- Support of current PSAAP collaborations and the development of PSAAP2 program
- Management of the SAIC contract to provide various administration support to HQ

Preliminary planned activities in FY13:

- Host external advisory panels
- Support programmatic needs of the DOE Exascale Initiative
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program

WBS 1.5.5.3 Applications in Support of Manufacturing Production and Connectivity (Y-12)

The Applications in Support of Manufacturing Production and Connectivity project supports the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. The project includes support

for connecting to ASC computing resources and job submission, execution, and visualization. The project provides the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources. Development and deployment of software to support the solution of manufacturing problems is also supported by the project. Visualization techniques that can be utilized in the Y-12 network and computing infrastructure will be evaluated and implemented. Finally, participation in Nuclear Weapons Complex ASC-related activities is covered.

In FY11, Y-12 installed and utilized the parallel version of Attila, a deterministic radiation transport code, for the determination of “adjoint” maps for nuclear facilities where fixed detectors are installed. Y-12 also gained expertise in use of the SIERRA tools through the collaboration efforts with KCP. Finally, Y-12 utilized MCNP software on the Y-12 cluster to simulate radiation transport in selected regions of the Y-12 Uranium Processing Facility in support of the design for the Criticality Accident Alarm System, which responds and alarms to a nuclear criticality.

Planned activities in FY12:

- Collaborate with the Kansas City Plant on determining physics-based simulation solution needs utilizing SIERRA framework
- Develop or utilize existing finite element models of secondaries and use SALINAS to estimate elastic or visco-elastic material properties in support of LEPs
- Evaluate GPU technology to support Monte Carlo and image processing applications
- Integrate an existing discrete-event nuclear facility simulation model and existing radiation transport models based on Attila to visualize the effects of nuclear excursions in time and space
- Utilize Y-12 and remote ASC cluster resources for production manufacturing problems

Expected deliverables in FY12:

- Determination of physics-based simulation needs at all production sites
- Demonstrated material modeling and manufacturing optimization codes on cluster
- Demonstrated discrete-event nuclear facility simulation model

Preliminary planned activities in FY13:

- Continue to apply new codes on the Y-12 cluster and to utilize Y-12 and remote ASC cluster resources for production manufacturing problems
- Continue evaluation of GPU computer as an alternative, general-purpose parallel computing resource
- Continue development of discrete-event nuclear facility simulation model
- Participate in Nuclear Weapons Complex ASC activities

WBS 1.5.5.4: System and Environment Administration and Operations

This level 4 product provides necessary operational support for reliable production computing and storage environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support.

WBS 1.5.5.4 System and Environment Administration and Operations (LLNL)

This product provides necessary operational support for reliable production computing environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support. Included within the scope of this product is the operational support for systems used as part of partnerships with academic, industrial, and other governmental agencies.

In FY11, LLNL deployed an unclassified GPGPU cluster (Edge) as well as a classified visualization cluster (Muir) in preparation for Sequoia. The Purple system was retired, including 17,000 classified disks. The Purple sub-floor infrastructure was completely cleared in preparation for Sequoia siting. Off-site disaster recovery backups now use tape encryption for increased security of data at rest and in transit.

Planned activities in FY12:

- Integrate TLCC2 systems and Sequoia
- Integrate TLCC2 and Sequoia parallel file systems
- Deliver Restricted Zone/Collaboration Zone system configurations
- Retire BlueGene/L, uBGL, Peloton class systems
- Self-maintain all compute platforms, including Sequoia and TLCC2
- Provide ongoing operational support: 24x7 system and facility monitoring support, DAT support, hardware repair
- Migrate core security infrastructure components to run on a RedHat Linux platform
- Develop user account and access management model that aligns with the architectural framework of the HPC Enclave
- Dispose of over 5000 drives and 26000 tapes (from retired STK Silos)

Expected deliverables in FY12:

- Sequoia and TLCC2 systems integrated

- BlueGene/L and Peloton-class systems retired
- HPC Enclave systems and network security model implemented
- Sequoia and TLCC2 GPFS integrated
- CRYPTOCard authentication service for HPC Enclave deployed
- RSA, Kerberos authentication, and LDAP directory service upgrades completed

Preliminary planned activities in FY13:

- Explore advanced security management and monitoring capabilities

WBS 1.5.5.4 System Administration and Storage (LANL)

The System Administration and Storage project covers all services for computational systems operated by LANL for the purpose of providing an HPC production computing environment for weapons designers, developers, and engineers. The project works with users to troubleshoot problems experienced while running their applications, and helps users transition from old to new computing platforms. The capabilities include system configuration, system and user security, resource management, system administration and monitoring, archival storage, Panasas, and NFS file systems.

In FY11, LANL enhanced production capability by accrediting all HPC clusters under a single, integrated security plan, and by deploying Cielo, Typhoon, Mapache, and ViewMaster for use by a multi-laboratory set of users.

Planned activities in FY12:

- Support system by conducting ongoing and daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Ensure workload is carried out by proper configuration of queues and scheduling policies; daily monitoring and problem resolution of user problems
- Ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Initially deploy the PLFS on generally available compute platforms for selected users

Expected deliverables in FY12:

- Support for Cielo weapons science and integrated weapons applications across the tri-lab community
- Contribution to the new releases of the CCE TOSS software stack, particularly support of the GPU programming environment for TLCC2
- Deployment of newer Panasas file system technology to replace aging technology and increase production capacity computing scratch file system capacity
- PLFS deployments on Cielo, Roadrunner, Redtail, Hurricane, Typhoon, Yellowrail, Turin, Mapache, and the TLCC2 clusters for selected users

Preliminary planned activities in FY13:

- Support a broader suite of weapons science and integrated weapons applications in capacity computing mode using GPUs

WBS 1.5.5.4 Operations and Procurement Support (LANL)

The Operations and Procurement Support project provides around-the-clock operations and monitoring of the scientific computing resources, including performance computers such as the original Roadrunner Base System (Redtail and Yellowrail), Roadrunner, Hurricane, Lobo, Typhoon, ViewMaster, Cielo, Mapache, and data storage and retrieval systems such as the HPSS. In addition to monitoring all components 24x7, the computer operators provide systems hardware maintenance for all ASC platforms. Working with the vendor system engineers, the operators also provide backup hardware support for the Cielo capability system. This includes all components of the production computing environment, including compute engines, hardware, file servers, archival storage systems, the facilities they reside in and utilities they are dependent upon, and all required software on these systems.

The procurement support aspect of this project assists customers with the technical and administrative aspects of planning, procurement, and contractual agreements for computer hardware and software products and services.

In FY11, LANL added the capability to perform hardware troubleshooting and repair for Cielo, Cielito and Appro hardware, including Hurricane, Typhoon, and Mapache. Additionally, this project developed formal maintenance and support policies, plus procedures for supporting the previously mentioned hardware.

Planned activities in FY12:

- Provide around-the-clock operations and monitoring of the scientific computing resources, including an increased level of system hardware self-maintenance for various computing and data storage systems and ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Provide hardware support for installation and integration of additional scientific computing platforms, including TLCC2
- Extend LANL and SNL best practices in systems operations, monitoring, and support via the ACES partnership
- Continue to provide technical and administrative support in the procurement of computer hardware and software products and services for all ASC platforms

Expected deliverables in FY12:

- Operations staff trained in the hardware repair of Cielo, Dell, and TLCC2 systems
- ACES best practices report
- Maintenance coverage for all ASC systems

Preliminary planned activities in FY13:

- Provide around-the-clock operations and monitoring of the scientific computing resources, including system hardware self-maintenance for all current and future generations of ASC platforms and data storage systems

WBS 1.5.5.4 Requirements Planning (LANL)

The Requirements Planning project collects and understands user requirements for production computing resources and quality of service, and to develop new metrics, data collection, and analysis techniques. The project is currently focused on establishing dependencies between cluster availability and cluster utilization and on implementing new technologies to accommodate rapidly growing information sets.

In FY11, the project deployed the common monitoring infrastructure, including new asset tracking software, to all computing clusters and most infrastructure elements (networks and file systems).

Planned activities in FY12:

- Upgrade the Zenoss HPC common monitoring software to the most recent base version of Zenoss to enable broad deployment to all infrastructure elements
- Enhance the Zenoss HPC interface for targeted troubleshooting
- Design a consolidated data management and reporting infrastructure to encompass system availability and utilization information coupled with environmental and application information

Expected deliverables in FY12:

- Upgraded version of Zenoss HPC deployed to all ASC clusters, including new capacity computing clusters, with data feeds from network switches and file systems
- Enhanced interface for troubleshooting delivered to HPC system administrators
- Design document for consolidated data management and reporting infrastructure

Preliminary planned activities in FY13:

- Develop and deploy consolidated data management and reporting infrastructure

WBS 1.5.5.4 Computing Platform Integration and Deployment (LANL)

The scope of the Computing Platform Integration and Deployment project is to accept delivery and begin deployment of production systems. Primary capabilities include completing the acceptance tests, diagnostics test, integrating the systems into the LANL yellow network, system stabilization, and transition into the classified network.

In FY11, the Cielo system was integrated into the LANL secure computing environment and was accredited for classified computing in February 2011, ahead of schedule.

Additionally, Cielo was upgraded to 1.37 petaFLOP/s in May. The CCC2 process started in July.

Planned activities in FY12:

- Coordinate with system vendor on the delivery of TLCC2 systems at LANL
- Develop acceptance test plan for TLCC2 systems at LANL
- Deploy initial FY12 TLCC2 cluster and provide system integration into the LANL secure and open networks
- Create plan for integrating TLCC2 into secure computing environment
- Develop integrated support structure with vendor for operational issues

Expected deliverables in FY12:

- Completed file system deployment on Cielo
- TLCC2 delivery, evaluation, diagnostic and acceptance testing, system integration into LANL network environment, system stabilization, security approval, and transition into classified network
- Deployment of TLCC2 for production workloads

Preliminary planned activities in FY13:

- Provide production support

WBS 1.5.5.4 Production Computing Services (SNL)

The Production Computing Services project's goals are to operate and maintain all ASC production platforms and associated support systems, and operate data services and visualization systems, long-term hierarchical storage services, high-performance network systems, tri-lab compatible cyber authentication and authorization systems, and monitoring and reporting services. This project supports tri-lab capability platform resource allocations and coordinates with tri-lab peers in establishing priority scheduling, if required. This project coordinates the integration and deployment of TLCC capacity systems into SNL's production computing environment, in collaboration with WBS 1.5.5.6 Common Computing Environment. Support of CCE common service and environment decisions and configuration management activities are also provided.

This project has expertise in operating capacity computing clusters; integrating file servers at the system or facility-wide level; deploying new computing, storage, and data management platforms; and in retiring end-of-life platforms. System administration for complex HPC environments is provided, as are design and development activities for new innovative computing platforms.

In FY11, SNL added the following new capabilities: support for national security computing needs (Cielo del Sur, Cray XE6 platform); prototype encrypted file system servers for Cielo del Sur; site planning and renovation for TLCC2 systems in Building 880, Annex, Room 230; decommissioned Black Rose and Red Rose systems to permit

realignment of Room 230 and improved cooling and power distribution. In addition, SNL acquired and installed an additional 3 petabytes of centralized Lustre file storage for TLCC2 systems, and acquired and installed Lynx platforms for data transfer to/from Cielo and LLNL.

Planned activities in FY12:

- Install and validate testing of TLCC2 systems delivered in FY12
- Assist in acquisition and installation/operation of next-generation test bed platforms
- Debug connectivity or user environment issues associated with access by LLNL and SNL personnel to Cielo platform at LANL
- Maintain and operate centralized Lustre file servers for unclassified and classified networks supporting TLCC1 and TLCC2 systems
- Maintain and operate centralized long-term hierarchical file storage systems running HPSS (classified and unclassified); upgrade to version 7.4
- Manage SNL computing resources via tri-lab EPR meetings, CCC process, and SNL Platform Oversight Committee process
- Demonstrate additional security capabilities within the NSCC to serve more than one customer group simultaneously
- Support the SNL FOUS Level 2 Milestone, *Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment*

Expected deliverables in FY12:

- Production ready TLCC2 platforms on restricted and classified networks
- Deployed next-generation computing test bed systems (as available)
- Production ready HPSS platform for national security environment
- Expanded file system production ready encrypted file system for NSCC systems
- Production environment for Cielo del Sur supporting national security customers
- Successful delivery of the SNL FOUS Level 2 Milestone, *Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment*

Preliminary planned activities in FY13:

- Continue operations as above for all capacity platforms, storage services, and national security systems
- Manage SNL computing resources via tri-lab EPR meetings, CCC process, and SNL Platform Oversight Committee process
- Integrate future platform into NSCC, either capacity or capability class
- Begin decommissioning plans for TLCC1 systems

WBS 1.5.5.5: Facilities, Network, and Power

This level 4 product provides necessary physical facility and other utility infrastructure. The following activities are included: facilities infrastructure, classified and unclassified facility networks, wide-area classified networks, ongoing network operations, infrastructure integration, and power.

WBS 1.5.5.5 Facilities, Network, and Power (LLNL)

The Facilities, Network, and Power project provides for the necessary physical facilities, utilities, and power capabilities to support staff and the ACE. Work in this area includes adequate raised floor space, cooling facilities, and power to site large-scale ASC platforms. In addition, this project funds needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff. Also included are classified and unclassified facility networks, wide-area classified networks, and ongoing network operations. This project also enables enhanced collaborations with academic and industrial partners.

In FY11, LLNL completed the Building 453 Direct Digital Control Upgrade project, as well as the final cell of the facility cooling tower. LLNL planned the Petascale Infrastructure project to support power and cooling efforts for Sequoia, which will be 85 percent complete by the end of the September 2011. LLNL deployed an InfiniBand SAN in the unclassified environment in preparation for Sequoia and a minimal redundant network capability in a separate building to enable failure recovery at a minimum level. LLNL delivered designs for the HPC Enclave network and the Inter-Site High Performance Computing (IHPC) network. LLNL continued to complete and implement energy savings initiatives resulting from self-benchmarking tools and computational fluid dynamic analysis for the TSF complex. Building 451 is now added to the list of LEED certified facilities, achieving LEED Silver in April 2011. LLNL also completed the Building 453 free cooling design to provide airside economization of the existing air handler system, which can result in a \$1M per year saving of electrical energy.

Planned activities in FY12:

- Track progress of the institution project elements for increased redundancy and reliability of the laboratory electrical distribution system that will support future petascale and exascale systems
- Continue analysis of future modifications and/or expansion of facilities that will be needed by future ASC systems
- Continue to implement the results from the self-benchmarking tool for all computer rooms in Building 453 (the TSF), Building 451, Building 439, Building 115, and Building 117; the self-benchmarking tool was created by DOE/SC and Lawrence Berkeley National Laboratory to continue to routinely identify prospective energy savings initiatives as computer rooms change¹⁰

¹⁰ See <http://hightech.lbl.gov/datacenters.html>.

- Continue to update computational fluid dynamics model for all ASC systems to profile the airflow required to cool the machines adequately as platforms are added and retired
- Clean out BlueGene/L and uBGL under-floor infrastructure
- Continue evaluation of 40GE and 100GE technologies
- Continue evaluation of QDR and EDR InfiniBand technologies
- Evaluate use of alternate vendor for 10G encryptor hardware

Expected deliverables in FY12:

- Completed installation of site preparation elements for Sequoia infrastructure installation, including the liquid cooling infrastructure, electrical distribution from bottom floor to the computer room, the electrical wiring and installation in the room itself, and mechanical and electrical infrastructure required to site Sequoia
- Upgrade of the Building 451 direct digital control
- Tools to perform power management of systems to prepare for exascale computing
- QDR InfiniBand and Ethernet networks for Sequoia and TLCC2 systems deployed
- HPC enclave network zones deployed

Preliminary planned activities in FY13:

- Continue to research the Building 453 Free Cooling project
- Continue to maintain and support the equipment in existing computational facilities
- Continue to analyze emerging network technologies

WBS 1.5.5.5 Facilities, Networking, and Power (LANL)

The Facilities, Networking, and Power project is responsible for the engineering, design, operation, and maintenance of the mission-important electrical, mechanical, cooling, network services, and other computing infrastructure in support of the ASC program. The project provides support for infrastructure design upgrades, project and space management, user interface and oversight, demolition and decommissioning of older systems, network backbones, user LANs, classified/unclassified network hardware and services, DisCom WAN, and computer site preparation for new platforms. Because the tri-lab community requires the systems to be operational at all times, the project provides on-call support after hours and on weekends for facility related issues.

In FY11, LANL implemented an automated system to measure and report PUE for the SCC. By using the automated system, LANL has reduced the PUE at the SCC from 1.65 to 1.32. Additionally, this project installed flow and temperature sensors on computer room air-conditioning units in the Laboratory Data Communications Center (LDCC). The measurements from these sensors will be used when calculating PUE for the LDCC facility. Finally, this project successfully integrated Cielo, ViewMaster, and Typhoon into the existing 12-lane network infrastructure.

Planned activities in FY12:

- Continue operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Decrease PUE and increase data center efficiency by implementing engineering strategies that will optimize energy use in the SCC and LDCC facilities
- TLCC2 site prep at SCC and LDCC computing centers
- Procure and install MARS software for enterprise sustainment modeling; populate MARS database with SCC facility data
- Enhance Zenoss HPC network monitoring features and redefine network to improve performance, reliability, and availability of backbone and interconnected services
- Implement out-of-band management for central networking infrastructure to improve manageability and performance

Expected deliverables in FY12:

- Substations integrated into exiting power distribution infrastructure at the SCC
- SCC and procurement plans for supporting pre-exascale and exascale systems based on ASC platform roadmap/timeline; plans will include timelines for required facility upgrades and procurements that will enable siting pre-scale and exascale systems at LANL
- MARS, a facility cost forecast system, deployed for SCC facility
- TLCC2 cluster integration into classified and unclassified networks
- Out-of-band management for HPC networking
- Zenoss HPC network monitoring features fully deployed

Preliminary planned activities in FY13:

- Implement FY13 portions of SCC and procurement plans developed in FY12
- Deploy MARS for LDCC and CCF facilities
- Continue operations and maintenance of electrical and mechanical systems and provide support for decommissioning of platforms that have reached end-of-life
- Decrease PUE and increase data center efficiency by implementing engineering strategies that will optimize energy use in the SCC and LDCC facilities
- Design follow-on to current 12-lane network infrastructure to accommodate future computing needs
- Investigate integration of within-cluster, short-term storage with cluster interconnect technologies

WBS 1.5.5.5 Facilities, Networking, and Power (SNL)

The Facilities, Networking, and Power project funds the power and space charges assigned to HPC systems (capacity and file system servers) and long-term hierarchical storage servers (running the HPSS software product). It provides for facilities and personnel to manage installation and removal of computing platforms, file systems, visualization systems, networking equipment, power distribution systems, and cooling systems in support of all computing resources. It also funds major operations contracts such as the ASC Distance Computing (DISCOM) WAN.

Facilities professionals have reduced overall operating expenses by minimizing cooling and electrical distribution expenses over the last several years through a comprehensive program of introducing more efficient computer room air conditioning units, using higher voltage electrical source power distribution units, exploring alternative energy sources and conservation mechanisms, which include reducing the volume of chilled water required for cooling and improving air flow in the facility by minimizing obstructions underneath the computer floor. These efforts have been recognized with several SNL-specific and national awards, including three 2011 EStar Awards from the DOE Office of Sustainability Support.

In FY11, SNL operations and research groups collaborated to install photoelectric cells, monitor their energy production, and credit the computing organization's electric charges with all energy generated by all cells in place at SNL's solar research locations. Experience gained through this activity will be used in future computing facility design decisions. Additional file system and tape storage capacity were added to the NSCC. Additionally, SNL successfully delivered reliable, quality service and access to any NNSA laboratory from any NNSA-designated ASC national user facility, independent of the location of the computing resource being utilized, supporting Multi-Site Target 3941.

Planned activities in FY12:

- Monitor performance of DisCom WAN and execute contract option for additional bandwidth if necessary
- Supply power and cooling as needed for all TLCC systems, file systems, and hierarchical storage systems
- Instantiate the IHPC network, as approved

Expected deliverables in FY12:

- Payment of power and space charges for all operational ASC systems
- Feasibility study and options for expansion of the NSCC
- Initial operational deployment of the IHPC network and limited resources, as approved by cyber security agreements

Preliminary planned activities in FY13:

- Create initial design for NSCC expansion or alternative facility location based on site-wide National Environmental Policy Act limitations

WBS 1.5.5.6: Common Computing Environment

The goal of the CCE product is to enable a common environment across the tri-labs that will initially be deployed on the TLCC systems. The scope of this product includes funded R&D projects to address gap areas identified by the tri-lab technical working groups.

The CCE working groups and projects focus on a common software stack, including but not be limited to, OS software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues such as configuration management, licenses, WAN access, and multi-realm security.

WBS 1.5.5.6 Tripod Operating System Software

TOSS is the tri-lab software stack that runs on Linux capacity clusters, initiating with TLCC platforms delivered in FY08. The goal of the TOSS project is to increase efficiencies in the ASC tri-lab community with respect to both the utility and the cost of the CCE. A notable achievement of this effort has been the porting of TOSS to SNL's RedSky system, given that its architecture more closely resembles many current capability systems rather than the typical commodity-based Linux capacity cluster.

This project delivers a fully functional cluster OS (kernel, Linux distribution, InfiniBand stack and related libraries, and resource manager) capable of running MPI jobs at scale on TLCC hardware. The system must meet CCE requirements for providing a common software environment on TLCC hardware across the tri-lab complex, now and into the future.

TOSS provides a complete product with full lifecycle support. Well-defined processes for release management, packaging, quality assurance testing, configuration management, and bug tracking are used to ensure a production-quality software environment can be deployed across the tri-lab in a consistent and manageable fashion.

In FY11, the TOSS project became more integrated across the tri-lab as the OS package matured. Highlights from this past year included release updates to TOSS (version 1.3-3, 1.3-4, and 1.4-2), which included security updates and bug fixes; development and deployment of a TOSS major update (version 1.4); deployment of TOSS on SNL's Red Sky class systems, requiring incorporation of needed Red Sky software modules into the core TOSS product; and continued development of TOSS 2.0, which is based on the next major RedHat release, RHEL 6.

Planned activities in FY12

- Provide ongoing TOSS software development and support
- Deploy TLCC2, including software integration and testing for the tri-lab environment
- Provide support for GPU-enhanced clusters used for visualization and GPGPU workloads
- Evaluate/incorporate new technologies already developed

Expected deliverables in FY12

- Development and deployment of GA release of TOSS 2.0, contingent on RedHat RHEL 6 schedule
- TLCC2 systems integration support provided by TOSS
- GPU-enhanced clusters integration support provided by TOSS

Preliminary planned activities in FY13

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 2.X (based on RHEL 6.X)

WBS 1.5.5.6 Open|SpeedShop

O|SS provides a wide range of performance experiments within a single environment. This includes support for PC sampling, inclusive and exclusive execution times for routines, hardware counters, as well as I/O and MPI tracing. O|SS provides all functionality through three fully interoperable interfaces, providing the greatest possible flexibility to users and code teams. Further, O|SS is designed to work on binaries of application without the need to recompile, enabling a clean and easy integration into the development workflow.

In FY11, the project continued deployment and support for ASC users and code teams. It continued the NNSA and DOE SC/ASCR co-funded project on O|SS, for which this CCE project provides the necessary funding. As part of this project, ASC started rewriting the O|SS framework for increased modularity and scalability and porting O|SS to existing NNSA capability machines, including Cielo and the Sequoia ID (Dawn), with the goal of ultimately providing the toolset on future NNSA systems, including Sequoia.

Planned activities in FY12:

- Provide an initial version of O|SS on Sequoia
- Integrate binary instrumentation
- Establish CBTF as the overall base for O|SS
- Complete the migration to the CBTF for all of O|SS
- Provide analysis and implementation of distributed aggregation functions

Expected deliverables in FY12:

- Production toolset running on the Sequoia ID (Dawn) and Cielo
- Initial prototype with limited capabilities for Sequoia
- Back port of all new additions for capability systems to TLCC clusters
- New hardware counter experiments following users' demands
- Initial Component Framework for scalable performance analysis constructed from existing O|SS components

Preliminary planned activities in FY13:

- Continue collaboration with OASCR on O|SS developments
- Complete implementation of new, scalable tool framework
- Provide additional analysis capabilities exploiting online aggregation

WBS 1.5.5.6 Workload Characterization

The WC Tool provides the capability to collect and report current and future requirements for compute resources (with programmatic characterization of the work) and computing resource usage (with programmatic characterization of the work). The tool includes a modular mechanism that can interface to multiple existing databases at each of the tri-labs. Additional development is needed in the areas of automated test suites and improved reporting capabilities.

Resource Management (SLURM/Moab) development is used to tie WC Tool demand/estimates, with their respective workload characterization, to job requests and resulting platform usage data.

In FY11, the project deployed the next major version (1.3) of WC Tool at all three labs, with improved demand/estimates interface and improved reporting interface for additional computing resource usage and machine utilization reports. The new version improved the tri-lab's ability to collect, analyze, and report compute resource requirements for current and future use, with programmatic characterization of the work. During FY11, the reports for platform demand and usage, with programmatic characterization of the work, continued to be available quarterly for the tri-lab. A new and custom reporting interface was developed to generate additional computing resource usage and machine utilization reports. This work included integration with the underlying platform usage databases (for example, the SLURM accounting database).

Planned activities in FY12:

- Enhance the WC Tool to improve the custom reporting interface and provide additional tools for configuration and administration
- Gather requirements from NNSA/HQ and each laboratory to expand the reporting capabilities
- Improve the WC Tool interface at each laboratory for data collection, storage, and reporting; track progress of work toward exascale systems for impact on workload characterization requirements
- Support the tri-lab Level 2 milestone, *Deployment of a Common Capacity Computing Environment*

Expected deliverables in FY12:

- Expanded WC Tool with improved custom reporting interface and additional reporting capabilities; additional configuration and administration tools

Preliminary planned activities in FY13:

- Assess usage and provide enhancements in tools or documentation and training
- Improve the WC Tool in the context of the tri-lab CCE projects
- Work with ASC HQ to meet new and/or expanded reporting requirements

WBS 1.5.5.6 SARAPE Redevelopment

SARAPE is a Web-based application that allows users within restricted domains to request selected CCE resources to which they are permitted access. A designated Guest Agent at the requestor's site provides verification and approval for each request. SARAPE manages each user's characteristics (work location, citizenship, clearance) and each resource's requirements (security level, access parameters) and tailors resource offerings to each user appropriately. A designated Host Agent at the resource site is responsible for final verification and approval of requests.

Phase 3 will be implemented in FY12. It addresses the APIs required to help interface SARAPE with other tools required to manage accounts among the tri-labs. As part of the IHPC deployment, a Service Catalog will be deployed through which collaborators can view and request accounts and services available in the shared environment. A proposal has been made to provide access to IHPC services by interfacing with Version 2 of SARAPE and additional tools, as identified to meet requirements of the IHPC. SARAPE version 2 will be capable of supporting many of the IHPC account management functions identified thus far.

In FY11, deficiencies in SARAPE version 1 were identified and communicated to the SARAPE team. Significant improvements and extensions to SARAPE were designed in a phased approach to address the majority of these deficiencies. Phase 1 and most of Phase 2 was completed in FY11 and included the addition of administrative functions at each lab.

Planned activities in FY12:

- Begin Phase 3 of SARAPE redevelopment
- Identify the account management requirements for IHPC services being implemented in FY12 (depends on IHPC use cases, ISA approval, and beginning implementation of initial services)
- Identify additional tools (commercial or otherwise) required to support IHPC services identified in use cases or being implemented

Expected deliverables in FY12:

- APIs interface with each lab local account services

Preliminary planned activities in FY13:

None planned.

WBS 1.5.5.6 Distributed Data Services

The DDS project seeks to define a common scheme for data organization and access such that the disparate and distributed information sources on ASC systems can be integrated for a system/environmental view for improved system analysis and to protect data access to appropriate roles. Such a system will allow cross-platform and cross-site development of mechanisms to take advantage of this integrated data resource, thus increasing the productivity of the platforms, developers, users, and administrators.

To date, application, OS, and hardware platform information have been collected with varying degrees of fidelity. The information gathered has largely been stored in log files that are only viewed, typically upon suspicion of error, either manually or by scripts tailored to look for particular indicators. This type of post run analysis and troubleshooting is very inefficient and will not scale as ASC systems continue to grow in both size and complexity. In addition, there are additional types of data (for example, file transfer activities and application run statistics) that could be integrated and be potentially used for active run-time participants, such as the applications, which could derive both increased performance and decreased debug time.

In FY11, work was done with each laboratory's cyber security departments to define and implement the security requirements for deploying this service on classified systems in the context of several use cases. Additionally, a commercial open source framework (WS02) was selected for building DDS.

Planned activities in FY12:

- Develop a production implementation of the FY11 prototype, including controls and a security model, that is compatible with the security constraints and regulations across SNL and LANL, for implementing a uniform data organization and access scheme to facilitate information sharing both within and across platforms but within a single HPC security domain
- Obtain security approval for use of ASC's production implementation of DDS in the classified environments, based on testing and validation first in the common test bed cluster environment and then on production clusters in respective unclassified environments

Expected deliverables in FY12:

- A security process established to enable DDS deployment; this work will be driven by a common set of test cases to be deployed on the unclassified networks and consistent with TOSS
- Document detailing system architecture and implementation plans, including controls and security model
- Catalog of data sources that address the FY12 use case deployments

Preliminary planned activities in FY13:

- Implement system in the classified environments according to FY11 design document and any modifications from FY12 deployments in the unclassified systems

- Deploy a prototype application on at least one system at each of the three labs that uses DDS services related to that system
- Work with developers, code maintainers, and system administrators to augment or adapt subsystems to better interoperate with DDS

WBS 1.5.5.6 Shared Workspace

The Shared Workspace project is the infrastructure for promoting collaboration across the laboratories. It currently includes the Gforge server that is housed and managed at SNL. Gforge administration support for the environment is done through identified staff at each laboratory. Gforge user licenses are purchased through LANL. During FY11, collaborative usage of Gforge by CCE working groups and projects increased for document management, maintaining wikis, managing subversion repositories, and tracking project activities.

Expectations for FY12 are to continue support for the Gforge server and licenses, and initiate additional services to support tri-lab account management to set the stage for additional cross-laboratory services.

In FY11, Shared Workspace provided a production collaboration environment for tri-lab CCE projects. CCE working groups and project teams now use Gforge for document management and other activities.

Planned activities in FY12:

- Maintain Gforge server for CCE project collaboration
- Integrate into the IHPC environment as a service
- Evaluate potential tools and decide on implementation methodology

Expected deliverables in FY12:

- Continued support for the Gforge server and licenses

Preliminary planned activities in FY13:

- Continue maintenance and increase use of Gforge and other Shared Workspace services

WBS 1.5.5.6 Debugger

All three labs are utilizing TotalView as the core debugger. LLNL has the more strategic relationship with TotalView Technologies and is also doing additional research into other debugging tools (for example, STAT). The approach to subset debugging is key as ASC moves toward larger scales and collects information that helps focus in on the trouble area. Input from all three labs and target applications are required. The working group would like to build a tri-lab debugger capability around the LLNL capability.

The focus of this project is to develop a more cohesive debug capability that is in line with the CCE goals of commonality and leverage of infrastructure and plans. This includes a range of sub-projects that include:

- Develop further the communication/support sub-team led through LLNL
- Host a common repository of debugger tool documentation on the CCE Gforge server
- Provide a central integration point for TotalView strategic planning through LLNL
- Provide cross-laboratory STAT deployment support (focused on subset debugging, which would also use TotalView and support scaling efforts; the University of New Mexico will provide guidance on porting and adapting the base infrastructure)
- Increase scalability and better integrate MRNet, which is the basis for STAT, through collaboration with the University of New Mexico
- Work towards addressing debugging on capability machines

In FY11, the project was able to deploy STAT on TLCC and non-TLCC machines (both SLURM and non-SLURM based) at each lab. In addition, milestones for the TotalView Scalability Project were met. Regarding MRNet, scalable process launching was implemented in a new layer called LIBI, and MRNet was ported to use this mechanism to launch its tree of processes rather than MRNet's internal rsh and ssh process launching.

Planned activities in FY12:

- Develop a working version of MRNet with the new interface and port STAT to it
- Continue coordination of issues and strategic needs for TotalView, and focus on GUI, memory tools, and infrastructure in the TotalView Scalability Project
- Develop the process and infrastructure to share documentation, training, issue resolution, and support
- Assess and develop plans and requirements for GPGPU debugger support

Expected deliverables in FY12:

- An implementation of STAT using the LIBI-based MRNet
- Increased debugger tri-lab support approach; further deployed TotalView and STAT
- Additional TotalView scalability completed based on contract

Preliminary planned activities in FY13:

- Increase the capability to scale debugging tools and to provide users with debugging tools that increase their productivity

WBS 1.5.5.6 Inter-Site High Performance Computing

With the advent of the TLCC effort, LANL, LLNL, and SNL have an increased need for integrated access to unclassified resources, which allows for improved collaboration and better utilization of computing resources. Attempts to facilitate such collaborations have

been hampered by incompatible security policies, interpretations, and implementations that make it difficult to access unclassified resources cross-site. These issues spawned the IHPC integration effort. This project has developed an architecture based upon a tri-labs ISA and a governance model. Tri-lab IHPC service deployment and network enhancements are the target for this year.

The target capability is to identify and implement a tri-lab security approach to achieve the following targets:

- Security policies and implementations that allow tri-lab access using identified protocols and technologies
- More effective resource control and utilization
- Utilization, where feasible, of home site security apparatus, including authentication equipment (for example, single sign on)
- User access to tri-lab resources within a specified period of time upon request and based on identified need
- Collaboration with the CCE user community to develop security policies that facilitate easier resource utilization
- Collaboration with the appropriate security personnel to ensure timely implementation
- Support of multiple classes of users, such as architecture and modeling, application, customer service, system administrators with root access, and end users
- Necessary network bandwidth and latency to cover current and projected requirements

In FY11, the ISA was signed by LLNL and the SNL designated approving authority. Service catalog and major usage scenarios were reviewed. Sites purchased the 10-Gbit/sec. encryptors and local security monitoring equipment in preparation for operation once tri-lab ISA and MOU signatures are completed.

Planned activities in FY12:

- Reestablish tri-lab cross-cell trust
- Develop detailed use case scenarios for IHPC to establish the bounds of IHPC operation (enclave versus desktop versus enterprise network) and how the agreement can best be leveraged to facilitate tri-lab communication/sharing in support of mission
- Upgrade to a 10-Gbit/sec. capability

Expected deliverables in FY12:

- Coordination of installation of 10-Gbit networking infrastructure based on ISA design
- Reestablishment of cross-cell trust

Preliminary planned activities in FY13:

- Leverage IHPC to implement and test disaster recovery mechanisms between tri-lab sites

WBS 1.5.5.6 Open Source Contract Maintenance

The Open Source Contract Maintenance project provides funding to outside developers who maintain tools and tool infrastructures that are critical for code teams or serve as the basis for internal tools. This funding will be provided to those developers through support contracts administered by LLNL, but each contract includes support for all three laboratories, and all three laboratories in close collaboration provide the technical guidance for the three contracts.

In FY11, the project continued with a structured process to provide support for some of the critical open source tools and tool components. During FY11, this project provided funding for O|SS, Valgrind, TAU, as well as supported workshops at three labs to increase visibility of ASC/CCE tools in the tri-lab user community.

Planned activities in FY12:

O|SS, Krell Institute

- Provide support in the form of support calls and code team interactions
- Provide training and tutorials at all three laboratories (one per laboratory)
- Expedite bug resolution and patch releases for the labs
- Perform continuous testing and release validation on laboratory systems

Valgrind, OpenWorks

- Provide support for tri-lab platforms (capacity and capability)
- Perform release validation on laboratory systems
- Expedite bug resolution and patch releases for the laboratories

TAU, Paratools

- Provide support in the form of support calls and code team interactions
- Provide training and tutorials at all three laboratories
- Coordinate among the tri-labs for contract execution

Expected deliverables in FY12:

- Support and maintenance for the three open source software components
- Quantifiable number of support hours delivered to the tri-labs
- Tutorials and/or workshops held at all three laboratories

Preliminary planned activities in FY13:

- Reevaluate which software components require support contract

- Continue support and maintenance for open source software components
- Coordinate among the tri-labs for contract execution

WBS 1.5.5.6 Open Message Passing Interface Integration/Scaling

The Open MPI Integration/Scaling project is targeted to develop a set of capabilities focusing on supporting scale increases, assessing performance of both MPI and user applications, and providing optimal parameters to users for better MPI performance. The complex multi-socket, multi-core NUMA node architecture of systems such as the TLCC mandates such an investigation. A close working relationship with the Open MPI community and other MPI developers is seen as a strategic need.

In FY11, the project provided both optimizations to Open MPI and cross MPI support for tools and scalability. For Open MPI, the project developed new shared memory backing facility framework to be used by Open MPI's run-time system for memory footprint reduction; implemented bug fixes to Open MPI for tri-lab issues, and developed memory registration optimizations on Linux systems that do not enable swap. These changes impacted TLCC machines and other IB systems. For MPI scalability support, the team tested prototype scaling test on tri-lab machines, developed the initial milestone with the MUST project and MPI Correctness tool; and used the Loba tool to collect and graph user application communication patterns.

Planned activities in FY12:

- Develop the MUST MPI Correctness tools, with ArgoNavis/Dresden
- Provide enhancements/bug fixes to Open MPI, ROMIO decoupling
- Modify and enhance Open MPI based on tri-lab need
- Share patches fed to MVAPICH community
- Customize user tools for application run parameters, assess MPI performance, identify best build and run parameters through studies, investigate impact of topology mapping, and provide results to end users
- Continue to improve MPI performance, reduce MPI memory footprint, and share MPI builds
- Continue to build tri-lab capability, scaling assessment tools, and performance enhancements geared toward tri-lab applications

Expected deliverables in FY12:

- Performance studies across different MPIs, Open MPI/MVAPICH/other MPI implementations
- Tools to support MPI at increased scale
- Modifications and enhancements to Open MPI, based on tri-lab need
- User tools to customize application run parameters

Preliminary planned activities in FY13:

- Continue to build tri-lab capability, scaling assessment tools, and performance enhancements geared toward tri-lab applications
- Focus the approach for implementation of mapping algorithms for applications and machine architecture

WBS 1.5.5.6 File System Architecture and Integration

The File System Architecture and Integration project integrates tri-lab efforts in developing I/O middleware to speed up I/O for codes that run on tri-lab and other HPC resources. This project leverages on-going efforts at LANL (PLFS) and LLNL (Scalable/Check-Point Restart Code (SCR)), and establishes both a strategy and initial implementations of a common standard API, that is, POSIX, to require no application changes, workout issues involved in a wider production installation and support process, and form the basis for working on a file systems architecture on which these and other projects can build. This supplements on-going efforts.

PLFS is a virtual parallel file system that reorganizes logical access to a single parallel file into physical access to multiple, non-shared files (that is, it converts N–1 parallel I/O into N–N). This has been demonstrated to improve checkpoint bandwidth by at least an order of magnitude on the three major parallel file systems: Lustre, GPFS, and PanFS.

SCR is an effort at LLNL with an identical motivation—to improve checkpoint bandwidth by eliminating file system contention. Additionally, SCR makes use of local storage devices by saving checkpoint data directly on the compute nodes' disk drives or to the nodes' memory.

In FY11, installations of both PLFS and SCR across the tri-labs were done for assessment and the initial prototype of an integrated capability was accomplished.

Planned activities in FY12:

- Continue tri-lab installments of PLFS/SCR and prototypes
- Continue development of architecture documents in file systems area
- Integrate PLFS capability further into the ADIO layer
- Integrate PLFS into a POSIX compliant mounted file system (via FUSE)
- Prototype a PLFS/SCR library that can be called directly from the application for increased performance/functionality
- Provide an MPI-IO interface for codes that use MPI-IO

Expected deliverables in FY12:

- Demonstration of an application writing an N–N checkpoint through PLFS, which internally uses SCR for storing its data and files but uses global storage for storing its metadata

- Demonstration of an application that writes each checkpoint as a single shared file using MPI-IO integration
- Demonstration of an application that writes each checkpoint as a single shared file using POSIX utilizing SCR via PLFS

Preliminary planned activities in FY13:

- Continue I/O integration strategy
- Coordinate among the tri-labs for I/O strategy/capability

WBS 1.5.5.6 General-Purpose Computing on Graphics Processing Units Programming Environment

The GPGPU Programming Environment project will assess the impact of GPGPU integration in HPC environments and collect user requirements for a production tri-lab GPGPU environment, including compilers, libraries, runtime environment, debugger, performance, and other tools. This initial environment will be deployed on production GPGPU clusters, such as TLCC2 GPGPU clusters at LANL, and user feedback collected. Based on this user feedback, refinements will be made and wherever gaps are identified, an approach for bridging these gaps in FY13 will be suggested.

Planned activities in FY12:

- Collect user requirements for production GPGPU environment
- Provide input to OJSS project for GPGPU future support
- Provide input to debugger project for GPGPU effort integration
- Collect and analyze user feedback regarding this environment

Expected deliverables in FY12:

- Initial GPGPU programming environment deployed
- Gap analysis based on user feedback completed

Preliminary planned activities in FY13:

- Continue assessing requirements for production GPGPU environment
- Deploy and configure GPGPU environment; collect user feedback

WBS 1.5.5.6 Resource Manager Capability Integration

The Resource Manager Capability Integration project will increase the support level for SLURM (resource manager that is part of the CCE software stack) across the three laboratories. SLURM development is moving primarily to a commercial entity. While this is the model that ASC strives for with products originating at the labs, it is necessary to ensure the tri-labs needs are met as SLURM grows in use by other organizations. This

project will integrate the support at each laboratory, including collection and coordination of requirement, and ensure that testing for the ASC workload and requirements is done.

Planned activities in FY12:

- Coordinate resource manager requirements
- Provide SLURM development support targeted to laboratory needs
- Initiate requirement development targeted toward future resource management needs, such as GPGPU and exascale

Expected deliverables in FY12:

- SLURM usage scenarios, testing suite, and development tasks targeting tri-lab needs
- Initial requirements study for future capability needs

Preliminary planned activities in FY13:

- Continue assessing requirements for future capability needs

V. ASC Level 1 and 2 Milestones

Table V-1. ASC Level 1 *Proposed* Milestones and Interfaces with Defense Programs Components from FY12–FY16

Milestone Title	Level	FY	Completion Date	Site(s)	Participating Program Offices
Develop, implement, and apply a suite of physics-based models and high-fidelity databases necessary for predictive simulation of the initial conditions for primary boost (initial conditions 1)	1	FY12	Q4	LANL, LLNL	Science Campaigns ASC Campaign
Assessment of weapon surety status	1	FY13	TBD	SNL	ASC Campaign Engineering Campaigns
Demonstrate predictive capability for weapon system response to short-pulsed neutrons in hostile radiation environment	1	FY13	TBD	SNL	ASC Campaign
Baseline demonstration of UQ aggregation methodology for full-system weapon performance prediction	1	FY14	TBD	LANL, LLNL, SNL	Science Campaigns ASC Campaign DSW
Full-system safety assessment	1	FY14	TBD	SNL	ASC Campaign Engineering Campaigns
Advanced models to support initial conditions for boost (initial conditions 2)	1	FY14	TBD	LANL, LLNL	Science Campaigns ASC Campaign

Table V-3. Quick Look: Level 2 Milestone Dependencies for FY12¹¹

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram	Site
TBD	Improved Physics Fidelity Relevant to Nuclear Performance and Aboveground Experiment Simulations	2	FY12	9/30/12	IC	LLNL
TBD	Enhanced Multi-Physics Arbitrary Lagrangian-Eulerian Capability	2	FY12	9/30/12	IC	LLNL
TBD	High-Pressure Plutonium Strength Model within a Baseline Continuum Model Framework	2	FY12	3/31/12	PEM	LLNL
TBD	Predictive Capability Assessment Project Comparison of Suite Calibrated Event Calculations to Annual Assessment Review Events	2	FY12	9/30/12	V&V	LLNL
TBD	Development and Application of an Improved Uncertainty Quantification Methodology for a Comprehensive Uncertainty Quantification Analysis Using the Secondary Computational Assessment Metrics Project Suite	2	FY12	9/30/12	V&V	LLNL
TBD	Lorenz Simulation Interface Beta Release	2	FY12	12/31/11	CSSE	LLNL
TBD	Early Users on Unclassified Sequoia Hardware	2	FY12	9/30/12	CSSE/FOUS	LLNL
TBD	High Performance Computing Enclave	2	FY12	9/30/12	FOUS	LLNL
TBD	Tri-Lab Linux Capacity Cluster-2 Early Classified Use	2	FY12	6/30/12	FOUS	LLNL
TBD	Improve Eulerian Application Codes through Additional Physics and Algorithms to Allow Validation via High Energy Density Experiments	2	FY12	9/30/12	IC	LANL
TBD	User Release of the Neutron Transport Monte Carlo Application Toolkit	2	FY12	9/20/12	IC	LANL
TBD	A Common Mix Model for Multi-Fluids	2	FY12	9/30/12	PEM	LANL
TBD	Benchmark Evaluation of Predictive Capability for Boost Using LANL Boost Validation Suite	2	FY12	6/30/12	V&V	LANL
TBD	Application Deployment of a Quick Parallel Log-Structured File System Capability	2	FY12	6/30/12	CSSE	LANL
TBD	Visualization Cluster Upgrade Project	2	FY12	12/31/12	CSSE	LANL
TBD	LANL Tri-Lab Linux Capacity Cluster-2 Production Readiness	2	FY12	6/30/12	FOUS	LANL

¹¹ Factors such as FY12 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

TBD	Particle Radiation Transport Predictive Capability to Enable Cavity System-Generated Electromagnetic Pulse Predictions in Realistic Re-Entry Body Geometry	2	FY12	9/30/12	IC	SNL
TBD	Large Eddy Simulation Capability for B61 Qualification Activities	2	FY12	6/30/12	IC	SNL
TBD	Integrated Workflow and Problem Setup for SIERRA	2	FY12	3/31/12	IC	SNL
TBD	Coupled Large Eddy Simulation and Structural Dynamics (SALINAS) Capabilities for B61 Captive Carry	2	FY12	9/30/12	PEM	SNL
TBD	Density Functional Theory Assessment of Defects in Neutron-Irradiated III-V Materials	2	FY12	9/30/12	PEM	SNL
TBD	Failure Predictions for B61 Abnormal Mechanical Environments	2	FY12	9/30/12	V&V	SNL
TBD	Intrinsic Verification and Validation Enabled by SIERRA/DAKOTA Integration	2	FY12	9/30/12	V&V	SNL
TBD	Demonstration of a Legacy Application's Path to Exascale	2	FY12	3/31/12	CSSE	SNL
TBD	Characterize the Role of the Mini-Application in Predicting Key Performance Characteristics of Real Applications	2	FY12	6/30/12	CSSE	SNL
TBD	Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment	2	FY12	6/30/12	FOUS	SNL
TBD	Cielo Production Capability Readiness	2	FY12	12/31/11	CSSE/FOUS	LANL, SNL
TBD	Deployment of a Common Capacity Computing Environment	2	FY12	9/30/12	FOUS	LLNL, LANL, SNL

Table V-4. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY13

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)
TBD	Application of Nuclear Performance Codes to Inertial Confinement Fusion Relevant Problems	2	FY13	9/30/13	IC	LLNL
TBD	Demonstrate Erosion with Auto-Contact Capability	2	FY13	12/31/12	IC	LLNL
TBD	Upgrade to Plutonium Equation of State that Benefits from Equation of State Variations Study	2	FY13	9/30/13	PEM	LLNL
TBD	Late-Time Studies of Curium and Americium Diagnostics	2	FY13	9/30/13	PEM	LLNL
TBD	Predictive Capability Assessment Project Comparison of Suite Calibrated Event Calculations, Annual Assessment Review Event Calculations, and Best Physics Calculations	2	FY13	9/30/13	V&V	LLNL
TBD	Early Users on Classified Sequoia Hardware	2	FY13	3/31/13	FOUS	LLNL
TBD	Implement in Lagrangian Application Codes the Physics and Infrastructure Improvements Required for the FY14 Predictive Capability Framework Pegpost for Initial Conditions 2	2	FY13	9/30/13	IC	LANL
TBD	User Release in the Eulerian Application Codes of an Alternative Hydrodynamics Option	2	FY13	6/30/13	IC	LANL
TBD	Development and Release of Improved Physics Required to Support FY14 Predictive Capability Framework Initial Conditions 2	2	FY13	6/30/13	PEM	LANL
TBD	Quantification of Uncertainty Due to Numerical Errors and Approximations in Multi-Physics Calculations	2	FY13	6/30/13	V&V	LANL
TBD	Programming Models and Data Analysis Environments for Extreme-Scale Systems	2	FY13	9/30/13	CSSE	LANL
TBD	HPSS v8.x Deployed in the Classified and Unclassified Networks	2	FY13	9/30/13	FOUS	LANL
TBD	Fully Functioning Integrated SIERRA/Solid Mechanics-Structural Dynamics Application	2	FY13	9/30/13	IC	SNL
TBD	Coupled High/Low Mach Number Algorithms	2	FY13	9/30/13	IC	SNL
TBD	High-Frequency Electromagnetic Field Capabilities in EIGER for Re-Entry Body/Re-Entry Vehicle Simulations	2	FY13	6/30/13	IC	SNL

TBD	Exascale Capability for Solvers and Algorithms Software Stack	2	FY13	9/30/13	IC	SNL
TBD	Physics-Based Model for Lightning Arrestor Connector Performance	2	FY13	9/30/13	PEM	SNL
TBD	Realistic Material Variability with Mechanical Constitutive Models for Directed Stockpile Work Materials	2	FY13	9/30/13	PEM	SNL
TBD	Computational Uncertainty Quantification for the QASPR Project	2	FY13	9/30/13	V&V	SNL
TBD	Intrinsic Verification and Validation Enabled by SIERRA/DAKOTA Integration	2	FY13	9/30/13	V&V	SNL
TBD	Data Co-Processing for Extreme Scale Analysis	2	FY13	12/31/12	CSSE	SNL
TBD	Peer-to-Peer File System Performance	2	FY13	3/31/13	CSSE	SNL
TBD	Deployment of a Common Capacity Computing Environment	2	FY13	9/30/13	FOUS	LLNL, LANL, SNL

Detailed Milestone Descriptions for FY12

Milestone (ID#): Lorenz Simulation Interface Beta Release		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: 12/31/11		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone builds upon the Lorenz job management application, a Web-based tool for submitting and managing batch jobs. FY12 work will encompass the following new Lorenz deliverables: 1) a Web-based interface for specifying and launching a vertical application of interest to ASC, and 2) the associated Web service infrastructure for supporting these operations. The Lorenz project is focused on making HPC easier and more efficient through the use of state-of-the-art Web technologies.		
Completion Criteria: This milestone is complete when users are able to use the Lorenz Web application to populate a simulation input deck, as well as launch and monitor jobs on both the classified and unclassified networks at LLNL.		
Customer: SSP and weapons science customers.		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The “handoff” of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: LC Division personnel and ASC customers.		
Supporting Milestones:		
Program	Title	Due Date
N/A		
Codes/Simulation Tools Employed: None		
Contribution to the ASC Program: Improved ability to manage, monitor, and launch simulation jobs running on ASC supercomputers.		
Contribution to Stockpile Stewardship: Improved ability to manage, monitor, and launch simulation jobs running on ASC supercomputers.		

Milestone (ID#): Lorenz Simulation Interface Beta Release				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Schedule slips	Low	Low	Low

Milestone (ID#): Early Users on Unclassified Sequoia Hardware				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: 9/31/12				
ASC nWBS Subprogram: CSSE/FOUS				
Participating Sites: LLNL				
Participating Programs/Campaigns: ASC				
Description: The Sequoia final system is planned to provide peak compute power of over 20 petaFLOP/s of computing cycles to the weapons program. The Sequoia system will be delivered in phases beginning in the first quarter of FY12, with final rack deliveries early in the third quarter of FY12. The goal of this milestone is to have early users successfully running their codes on some portion of the Sequoia system on the unclassified network.				
Completion Criteria: Racks have been assembled in the TSF and have run at least one user's ported code. A user will write a memo certifying that their code has run on Sequoia.				
Customer: ASC				
Milestone Certification Method: A report will be written as a record of milestone completion. The memo from the user certifying that his/her code has run will also be submitted.				
Supporting Resources: ASC program people, the Sequoia system, IBM				
Supporting Milestones:				
Program	Title	Due Date		
N/A				
Codes/Simulation Tools Employed: Sequoia software environment				
Contribution to the ASC Program: Advanced architecture resource for UQ, 20 petaFLOP/s peak capability				
Contribution to Stockpile Stewardship: Acceleration of ability to do UQ.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Hardware issues	Moderate	Low	Low

Milestone (ID#): High Performance Computing Enclave				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: 9/30/12				
ASC nWBS Subprogram: FOUS				
Participating Sites: LLNL				
Participating Programs/Campaigns: ASC				
Description: The HPC Enclave is a security environment – a restricted zone and a collaboration zone that segregates compute platforms such that the collaboration zone allows foreign nationals access to LC resources while disabling access to restricted zone resources. Implementing the enclave requires completing a security plan, extensive network and system configuration changes and administrative controls.				
Completion Criteria: This milestone is complete when 1) the security plan is complete and 2) users are running on platforms in both the restricted and collaboration zones while sharing file system and archive resources.				
Customer: ASC				
Milestone Certification Method: A report will be written as a record of milestone completion. The security plan documentation will also be submitted.				
Supporting Resources: CSP, ASC				
Supporting Milestones:				
Program	Title	Due Date		
N/A				
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: Ability to meet demands of multiple user sets while minimizing costs of storage resources.				
Contribution to Stockpile Stewardship: Increase accessibility of clusters while maintaining security and containing costs.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Hardware issues	Moderate	Low	Low
2.	Security issues	Moderate	Low	Low

Milestone (ID#): Tri-Lab Linux Capacity Cluster-2 Early Classified Use				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: 6/30/12				
ASC nWBS Subprogram: FOUS				
Participating Sites: LLNL				
Participating Programs/Campaigns: ASC				
Description: The TLCC2 classified system named Zin is an 18-SU cluster of Intel Sandybridge technology with Qlogic QDR InfiniBand interconnect. It will provide approximately 850 peak teraFLOP/s of capacity compute cycles to the center. The system is scheduled to be delivered in phases, beginning in October 2011, with delivery of all racks completed by December 2011. The goal of this milestone is to have early users running their codes on Zin on the classified network.				
Completion Criteria: Racks have been assembled in the TSF, the system has been moved to the classified network, and early users have run their codes. A user will write a memo certifying that their code has run on Zin.				
Customer: ASC				
Milestone Certification Method: A report will be written as a record of milestone completion. The memo from the user certifying that his/her code has run will also be submitted.				
Supporting Resources: ASC program people, TLCC2 vendors				
Supporting Milestones:				
Program	Title	Due Date		
N/A				
Codes/Simulation Tools Employed: TLCC2 software environment.				
Contribution to the ASC Program: Capacity cycles needed for stockpile stewardship.				
Contribution to Stockpile Stewardship: Capacity cycles needed for stockpile stewardship.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Vendor hardware or software defect delays integration schedule	Moderate	Moderate	Low

Milestone (ID#): Application Deployment of a Quick Parallel Log-Structured File System Capability		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: 6/30/12		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
<p>Description: Current petascale supercomputers have been carefully provisioned to provide a computing platform well balanced for current workloads. Unfortunately, the I/O patterns of many of these workloads are poorly matched to the underlying storage systems; for these workloads, achieving even a fraction of aggregate spindle speed can be very challenging. PLFS is a virtual file system that rearranges workload I/O into patterns better matched to the underlying storage system. PLFS addresses issues with two different workloads. Workloads where all processes write into a shared file (N to 1) with small, unaligned writes performing badly on all parallel file system. PLFS rearranges this workload into non-shared file access pattern similar to each process writing its own file, which performs much better for parallel data operations. Workloads where each process writes its own file (N to N) suffer from poor initial creation due to simultaneous inserts into a single directory. PLFS addresses this file creation into a shared directory by hashing creates over multiple directories. PLFS provides both of these traffic-shaping functions for these two dominant workloads with no substantive changes to applications.</p>		
<p>Completion Criteria: Demonstration of a production deployed PLFS using a production weapons or weapons science code that uses the N to 1 I/O pattern for checkpoint and restart, and shows substantive speed up for I/O times compared to using the scratch parallel file system directly on a to-be-determined scale on an ASC machine.</p> <p>Demonstration of either an N to 1 PLFS that has been converted to N to N or a direct N to N style production weapons or weapons science code on a to-be-determined scale on an ASC machine showing checkpoint speed ups over direct scratch file system use due to hashing of file creates. Problem sizes/scales, speed-ups, PLFS production version, codes used, and code versions used will be documented.</p>		
Customer: Integrated code developers and users on ASC production platforms.		
<p>Milestone Certification Method:</p> <p>A program review is conducted and its results are documented.</p> <p>Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.</p>		
Supporting Resources: System resources and participation of code teams		

Milestone (ID#): Application Deployment of a Quick Parallel Log-Structured File System Capability				
Supporting Milestones:				
Program	Title	Due Date		
ASC/CSSE	Cielo capability computing platform production capability readiness	June 2011		
Codes/Simulation Tools Employed: PLFS file system and IC (for example, EAP codes).				
Contribution to the ASC Program: Improves throughput of production codes through significantly enhanced I/O performance and boosts productivity for code developers in effectively exploiting high-performance file systems.				
Contribution to Stockpile Stewardship: Faster time-to-solution for large-scale simulations.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Loss of key personnel	High	Moderate	Medium
2.	Schedule conflicts with production platform needs	Moderate	Moderate	Medium

Milestone (ID#): Visualization Cluster Upgrade Project				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: 12/31/2011				
ASC nWBS Subprogram: CSSE				
Participating Sites: LANL				
Participating Programs/Campaigns: ASC				
Description: This milestone provides the necessary resources to visualize output generated on petascale clusters. Requirements will be developed, equipment will be purchased, and the cluster will be integrated into the LANL computer center. The milestone will be complete when the cluster is ready for production work.				
Completion Criteria: Results documented as a report. System made available to the ASC community.				
Customer: Weapons user community.				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: Access to Roadrunner and other supercomputer resources.				
Supporting Milestones:				
Program		Title		Due Date
ASC/CSSE		Cielo capability computing platform production capability readiness		June 2011
Codes/Simulation Tools Employed: ASC code projects and graphics tools to demonstrate readiness				
Contribution to the ASC Program: The milestone will provide the ASC community with the ability to visualize large-scale simulation datasets coming from the petascale platforms of Roadrunner and Cielo, by acquiring and standing up a visualization compute platform dedicated to this task.				
Contribution to Stockpile Stewardship: The milestone will provide the ability to visualize large-scale simulation datasets relevant to stockpile stewardship.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Loss of key personnel	High	Moderate	Medium
2.	Unavailability of selected high-end graphics cards in the milestone time frame	High	Low	Medium

Milestone (ID#): Visualization Cluster Upgrade Project				
3.	Key personnel unavailable because committed to other milestones or high-priority activities	High	Moderate	Medium
4.	Difficulty in integrating platform because of issues with new technology introduced	High	Moderate	Medium

Milestone (ID#): LANL Tri-Lab Linux Capacity Cluster-2 Production Readiness				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: 6/30/12				
ASC nWBS Subprogram: FOUS				
Participating Sites: LANL				
Participating Programs/Campaigns: ASC				
Description: The TLCC2 platform is ready for capacity production work at LANL. The requirements include: machine accessibility and integration, operational support, user testing of applications has been done, system reliability has been achieved, application and I/O performance testing has been demonstrated, and a milestone review has been completed.				
Completion Criteria: TLCC2 is ready for production. The system has demonstrated an acceptable production user environment with all the associated support, testing, reliability, and applications use of the system.				
Customer: NNSA/ASC HQ, SSP, LANL weapons applications community.				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. This will follow from the program review and will be submitted to HQ, including the report and all the review presentations.				
Supporting Resources: CSSE, FOUS, platform funding, LANL facilities.				
Supporting Milestones:				
Program	Title	Due Date		
N/A				
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: Provides production capacity compute cycles to ASC Program, including scalable performance.				
Contribution to Stockpile Stewardship: Production capacity platform for the ASC Program.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	System not ready for production	High	Low	Medium

Milestone (ID#): Demonstration of a Legacy Application's Path to Exascale		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: 3/31/12		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
<p>Description: Cielo is expected to be the last capability system on which existing ASC codes can run without significant modifications. This assertion will be tested to determine where the breaking point is for an existing highly scalable application. The goal is to stretch the performance boundaries of the application by applying recent CSSE R&D in areas such as resilience, power, I/O, visualization services, SMARTMAP, lightweight LWKs, virtualization, simulation, and feedback loops. Dedicated system time reservations and/or CCC allocations will be used to quantify the impact of system-level changes to extend the life and performance of the ASC code base. Finally, a simulation of anticipated exascale-class hardware will be performed using SST to supplement the calculations.</p>		
<p>Completion Criteria: Completion of program review and final document published as a SAND report.</p>		
Customer: ASC IC		
<p>Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.</p>		
Supporting Resources: Cielo, and associated test beds		
Supporting Milestones:		
Program	Title	Due Date
N/A	N/A	N/A
Codes/Simulation Tools Employed: Charon, CTH, and associated mini-apps		
<p>Contribution to the ASC Program: Supercomputing systems are evolving to new architectures. In order for ASC's investment in existing codes to remain viable, it is important to understand and quantify how far they can be extended before dramatic source code modifications are required. The CSSE-developed system software being exercised by this milestone provides performance enhancements that make use of new hardware technologies without requiring changes to the application.</p>		
Contribution to Stockpile Stewardship: More effective use of capability platforms		

Milestone (ID#): Demonstration of a Legacy Application's Path to Exascale				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Insufficient access to capability platforms—particularly in dedicated mode	High	Low	Low

Milestone (ID#): Characterize the Role of the Mini-Application in Predicting Key Performance Characteristics of Real Applications		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: 6/30/12		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: The Mantevo project includes a set of application proxies, referred to as mini-apps, and designed by code developers to represent key runtime performance characteristics of their applications. SNL will analyze two of these mini-apps to determine how well they represent the full application programs. Specifically, SNL will profile the runtime performance of the mini-app and application, characterizing the relationship between the two on at least two HPC platforms (including Cielo).		
Completion Criteria: Completion of program review and final document published as a SAND report.		
Customer: ASC IC		
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.		
Supporting Resources: Cielo and associated test beds		
Supporting Milestones:		
Program	Title	Due Date
N/A	N/A	N/A
Codes/Simulation Tools Employed: Mantevo mini-apps and represented codes.		
Contribution to the ASC Program: Guidance on addressing issues impacting key performance areas in the full applications.		
Contribution to Stockpile Stewardship: The long-term usefulness of ASC codes to the SSP will necessitate that the codes evolve as supercomputer architectures change. Dramatic changes to full production codes are expensive and high risk. Applications proxies can play an important role in determining the evolutionary path of production codes without incurring the high overhead of working with the full code.		

Milestone (ID#): Characterize the Role of the Mini-Application in Predicting Key Performance Characteristics of Real Applications				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Insufficient profiling data	High	Low	Low
2.	Insufficient access to large scale	High	Low	Low

Milestone (ID#): Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: 6/30/12		
ASC nWBS Subprogram: FOUS		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: The ASC program is acquiring a second generation of TLCC2 beginning in Q4 FY11. All equipment will be delivered by the end of Q2 FY12. This milestone represents integration and deployment of the clusters delivered to SNL onto the SRN and SCN networks for production use.		
Completion Criteria: The platforms will be operational, stable, and running the standard TOSS user environment at full scale. File system will be operating at sufficient scale and performance to support SNL workload. ASC and DSW programs are running and generating correct results.		
Customer: NW Program		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The “handoff” of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: Tri-lab CCE TOSS project and SNL FOUS operations teams		
Supporting Milestones:		
Program	Title	Due Date
CSSE	TLCC2 Contract Awarded	FY11Q2
CSSE	Deploy TLCC2 and a common capacity computing environment	FY12Q4
Codes/Simulation Tools Employed: TOSS common operating environment		
Contribution to the ASC Program: Efficient operations and CCE.		
Contribution to Stockpile Stewardship: Additional capacity to support B-61 LEP, V&V and QMU calculations.		

Milestone (ID#): Integrate the Tri-Lab Linux Capacity Cluster-2 Clusters Delivered in FY12 into SNL Environment				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Delay in delivery of platforms	Medium	Low	Low
2.	Hardware error discovered late in procurement/delivery schedule	High	Low	Medium

Milestone (ID#): Cielo Production Capability Readiness								
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC						
Completion Date: 12/31/11								
ASC nWBS Subprogram: CSSE								
Participating Sites: LANL, SNL								
Participating Programs/Campaigns: ASC								
<p>Description: This milestone is based on meeting the requirements as outlined in the Production Capability Readiness Milestone Description Guidelines. These requirements certify, through a review, that Cielo is ready for production capability. These requirements include: machine accessibility and integration, operational support, usage model has been demonstrated, user testing of applications has been done, system reliability has been achieved, application and I/O performance testing has been demonstrated, and a milestone review has been completed.</p>								
<p>Completion Criteria: Follows the ASC Level 2 Milestone criteria for capability platforms: that all of the topics identified in the description above have been successfully demonstrated for capability class simulations. These requirements are specifically listed in the usage model for Cielo, which defines that the system has demonstrated an acceptable production user environment with all the associated support, testing, reliability and applications use of the system.</p>								
<p>Customer: NNSA/ASC HQ, tri-lab ASC program managers responsible for CCCs, SSP, tri-lab weapons applications community.</p>								
<p>Milestone Certification Method:</p> <p>A program review is conducted, requiring all three laboratories to participate in the completion of the review. The review panel will include a mix of systems personnel and end users.</p> <p>Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. This will follow from the program review and will be submitted to HQ, including the report and all the review presentations.</p>								
<p>Supporting Resources: CSSE, FOUS, platform funding, ACES program managers, ACES staff, LANL facilities.</p>								
<p>Supporting Milestones:</p> <table border="1"> <thead> <tr> <th>Program</th> <th>Title</th> <th>Due Date</th> </tr> </thead> <tbody> <tr> <td>N/A</td> <td>N/A</td> <td>N/A</td> </tr> </tbody> </table>			Program	Title	Due Date	N/A	N/A	N/A
Program	Title	Due Date						
N/A	N/A	N/A						
Codes/Simulation Tools Employed: N/A								
<p>Contribution to the ASC Program: Provides production capability compute cycles to ASC Program including scalable performance.</p>								

Milestone (ID#): Cielo Production Capability Readiness				
Contribution to Stockpile Stewardship: Primary production capability platform for the ASC Program.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	System not ready for production	High	Low	Medium
2.	Files system reliability issues	High	Low	Medium
3.	File system performance not achieved	High	Low	Medium
4.	Applications do not scale well on the system	High	Low	Medium

Milestone (ID#): Deployment of a Common Capacity Computing Environment				
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC		
Completion Date: September 30, 2012				
ASC nWBS Subprogram: FOUS				
Participating Sites: LLNL, LANL, SNL				
Participating Programs/Campaigns: ASC				
Description: Deploy additional CCE capabilities for capacity computing environment, working towards a responsive and more efficient infrastructure to support computing for QMU and predictivity.				
Completion Criteria: Deploy CCE capabilities developed during FY11 including major upgrades to the common OS (TOSS 2.0); OJSS scalable framework; workload characterization; deployment of TotalView Scalability project work; and new capabilities in the security integration and OpenMPI projects. Demonstrate the tri-lab CCE software stack on the production ASC TLCC systems. Develop new capabilities in the continuing FY12 CCE projects. Prepare for deployment of the next generation of the ASC TLCC systems, which may include hardware and software integration and testing for the tri-lab environment. The tri-labs will continue to do gap and risk analysis of the CCE software stack and add new projects, as needed, to address high-priority gaps.				
Customer: Users of the ASC capacity systems				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: ASC FOUS sub-program				
Supporting Milestones:				
Program	Title		Due Date	
N/A	N/A		N/A	
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: CCE for capacity computing				
Contribution to Stockpile Stewardship: Easy-to-use capacity computing resource				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Resource availability	Low	Low	Low

Milestone Descriptions for Preliminary FY13

Milestone (ID#): Early Users on Classified Sequoia Hardware		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 3/31/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The 20 petaFLOP/s Sequoia system will be run unclassified for all 2012. In calendar year 2013, the system will be moved to the classified network and will run classified codes.		

Milestone (ID#): Programming Models and Data Analysis Environments for Extreme-Scale Systems		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: Emerging computer architectures pose new challenges for scientific applications at extreme scales. This milestone will deliver 1) programming models and supporting tools that allow for effective expression of large-scale concurrency while providing flexibility in mapping to the underlying architecture, and 2) a data analysis environment that addresses the challenges of massive data, including <i>in situ</i> analysis and advanced data reduction techniques. These products will be a step towards enabling productivity and performance at exascale.		

Milestone (ID#): HPSS v8.x Deployed in the Classified and Unclassified Networks		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: Implement HPSS version 8.x into the open and secure computing environments. This version provides the ability to run multiple core components of HPSS thus enabling archive scaling of meta-data processing.		

Milestone (ID#): Data Co-Processing for Extreme Scale Analysis		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 12/31/12		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
Description: ASC calculations produce complex datasets that are increasingly difficult to explore and understand using traditional post-processing workflows. To advance understanding of underlying physics, uncertainties, and results of ASC codes, SNL must gather as much relevant data as possible from large simulations. This drives SNL to couple data analysis and visualization capability with a running simulation, so that high fidelity data can be extracted and written to disk. This milestone deals with two alternate methods for analyzing large-scale data: <ul style="list-style-type: none"> • <i>In-situ</i> processing. One solution to providing analysis capabilities is to directly couple analysis libraries with the simulation. SNL has collaborated on developing the ParaView <i>in-situ</i> library, which is designed for this purpose. • In-transit processing. Another option for processing data at extreme scale is to loosely couple the running simulation with parallel data services that use separate processing resources to reliably perform a range of operations on the simulation data. Within this milestone, SNL will engineer, test, and evaluate a set of customer-driven data operations on large-scale data created by a running simulation. The data operations will be performed by instrumented versions of both the <i>in-situ</i> and in-transit solutions, with the resulting performance data published and made available to the ASC community.		

Milestone (ID#): Peer-to-Peer File System Performance		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 3/31/13		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC		
<p>Description: Exascale computing will likely require fundamental changes in the storage and management of persistent data. Incremental advances in current capabilities are likely inadequate. This milestone will provide performance analysis of a revolutionary approach to persistent storage—one that uses smart storage servers with access to a variety of different local and remote media (for example, disk, NVRAM, memory, and tape) and are pervasive throughout the computing platform. Storage servers have the ability to directly handle I/O requests, initiate third party transfers, or replicate the data as needed. Results will come from implementations for the Cielo system.</p>		

Milestone (ID#): Deployment of a Common Capacity Computing Environment		
Level: 2	Fiscal Year: FY13	DOE Area/Campaign: ASC
Completion Date: 9/30/13		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
<p>Description: Deploy CCE capabilities developed from FY12, including the next major release of common operating system and software stack. Continued deployment of the TLCC2 systems, which will include hardware and software integration and testing for the tri-lab environment.</p>		

VI. ASC Performance-Based Initiatives/Multi-Site Target Drivers for FY12–FY13

There is one Performance-Based Initiative target in the “Getting it Done Job List,” also known as the GTJDL, and the Multi-Site Target:

2.2 Achieve advances in experimental and computational tools used in resolving SFIs and in supporting LEP activities associated with early-phase primary implosion.

Implementing Criteria:

2.2.1 Refine experimental and computational tools that could enable the assessment of a future SFI.

Completion Criteria:

2.2.2 Meet the completion criteria for the associated Level 1 milestone for initial boost conditions, including pre-shot predictions for the Pollux experiment.

Participating sites: SNL, LANL, LLNL, NNSS

VII. Performance Measures

Table VII-1. ASC Campaign Annual Performance Results (R) and Targets (T)

REDUCED RELIANCE ON CALIBRATION: The cumulative percentage reduction in the use of calibration "knobs" to successfully simulate nuclear weapons performance (Long-term Outcome)

Reporting process: Labs will assess and report within two weeks of the end of Q3, progress toward Level-1 milestones for knob replacement and supporting Level-2 milestones. A briefing will be provided to the Predictive Science Panel at each meeting as a report on progress toward replacing the knobs. The PSP may provide advisory comments regarding the reported progress, but will not evaluate the progress.

At LLNL, progress on this metric is measured by completion of the N (i.e., 4) level-1 milestones that directly address knob replacement. Completing one knob-replacement level-1 milestone accumulates 100%/N (i.e., 100%/4=25%) reduced reliance on calibration. Partial progress toward the next knob-replacement level-1 milestones is achieved by successful completion of predictive capability level-2 milestones.

LANL: Progress on this metric will be reported 9/22/11
replacement Level 1 milestone accumulated by successful completion of predictive capability Level 2 milestones. Cumulative percentage is reported for July 1, 2005 through June 30, 2011 (FY05Q4 through FY11Q3). In 2011 we rebaselined this metric over 8 major pegposts of the Predictive Capability Framework, 2 pegposts per knob (N=8).

Evidence (Type): NA-10 Milestone Reporting Tool (MRT) reports (Original Documents)

	Q3 FY07	Q3 FY08	Q3 FY09	Q3 FY10	Q3 FY11	Q3 FY12	Q3 FY13	Q3 FY14	Q3 FY15	Q3 FY16	Q3 FY17
TARGETS	8%	16%	25%	30%	35%	40%	45%	50%	55%	60%	65%
LLNL	12%	18%	24%	33%	35%						
LANL	12%	18%	25%	33%	35%						
AVERAGE	12%	18%	25%	33%	35%						

VIII. Budget

Product / Project WBS	Lab	PEOPLE (\$Ms)	INFOSTRUCTURE (\$Ms)	CONTRACTS (\$Ms)	Total (\$Ms)
1.5.1 Integrated Codes					
Engineering and Physics Integrated Codes		85.914	0.000	0.000	85.914
	LLNL	31.037			31.037
	LANL	28.977			28.977
	SNL	25.900			25.900
	Other				0.000
Specialized Codes and Libraries		24.627	0.000	0.000	24.627
	LLNL	10.999			10.999
	LANL	6.328			6.328
	SNL	7.300			7.300
	Other				0.000
Applications and Algorithms Research		18.739	0.000	0.000	18.739
	LLNL	6.230			6.230
	LANL	8.009			8.009
	SNL	4.000			4.000
	Other	0.500			0.500
Applications Research for Next-Gen Platforms		15.627	0.000	20.500	36.127
	LLNL	5.427			5.427
	LANL	3.600			3.600
	SNL	3.600			3.600
	Other	3.000		20.500	23.500

Product / Project WBS	Lab	PEOPLE (\$Ms)	INFOSTRUCTURE (\$Ms)	CONTRACTS (\$Ms)	Total (\$Ms)
1.5.2 Physics and Engineering Models					
High Explosive		6.024	0.000	0.000	6.024
	LLNL	2.522			2.522
	LANL	3.502			3.502
	SNL				0.000
	Other				0.000
Equation of State		7.650	0.000	0.000	7.650
	LLNL	3.750			3.750
	LANL	3.900			3.900
	SNL				0.000
	Other				0.000
Nuclear Properties		5.900	0.000	0.000	5.900
	LLNL	2.000			2.000
	LANL	3.900			3.900
	SNL				0.000
	Other				0.000
Plasma and Radiative Properties		4.900	0.000	0.000	4.900
	LLNL	2.000			2.000
	LANL	2.900			2.900
	SNL				0.000
	Other				0.000
Advanced Hydrodynamics		4.900	0.000	0.000	4.900
	LLNL	1.500			1.500
	LANL	3.400			3.400
	SNL				0.000
	Other				0.000
Material Strength and Damage		13.615	0.000	0.000	13.615
	LLNL	3.635			3.635
	LANL	4.685			4.685
	SNL	4.950			4.950
	Other	0.345			0.345
Forensics and Cross-cutting Initiatives		11.240	0.000	0.000	11.240
	LLNL	4.352			4.352
	LANL	5.134			5.134
	SNL	1.754			1.754
	Other				0.000
Thermal and Fluid Response		5.100	0.000	0.000	5.100
	LLNL				0.000
	LANL				0.000
	SNL	5.100			5.100
	Other				0.000
Aerodynamics and Vibration		3.950	0.000	0.000	3.950
	LLNL				0.000
	LANL				0.000
	SNL	3.950			3.950
	Other				0.000
Radiation and Electrical Response		4.403	0.000	0.000	4.403
	LLNL				0.000
	LANL				0.000
	SNL	4.403			4.403
	Other				0.000
Russian Programs		2.250	0.000	0.000	2.250
	LLNL	0.750			0.750
	LANL	0.750			0.750
	SNL	0.750			0.750
	Other				0.000

Product / Project WBS	Lab	PEOPLE (\$Ms)	INFOSTRUCTURE (\$Ms)	CONTRACTS (\$Ms)	Total (\$Ms)
1.5.3 Verification and Validation					
V&V Methods		23.414	0.000	0.000	23.414
	LLNL	7.800			7.800
	LANL	10.214			10.214
	SNL	5.400			5.400
	Other				0.000
V&V Assessments		29.467	0.000	0.000	29.467
	LLNL	7.915			7.915
	LANL	5.852			5.852
	SNL	15.700			15.700
	Other				0.000
Data Validation, Archiving, SQA and Training		4.351	0.000	0.000	4.351
	LLNL	1.433			1.433
	LANL	1.568			1.568
	SNL	1.350			1.350
	Other				0.000

Product / Project WBS	Lab	PEOPLE (\$Ms)	INFOSTRUCTURE (\$Ms)	CONTRACTS (\$Ms)	Total (\$Ms)
1.5.4 Computational Systems and Software Environment					
Production Systems		2.179	29.200	0.000	31.379
	LLNL	1.092	16.300		17.392
	LANL	1.087	2.900		3.987
	SNL				0.000
	Other		10.000		10.000
Advanced Systems		3.900	85.915	6.000	95.815
	LLNL	0.100	85.915		86.015
	LANL				0.000
	SNL	3.800		3.000	6.800
	Other			3.000	3.000
System Software and Tools		20.602	0.000	0.000	20.602
	LLNL	6.195			6.195
	LANL	7.007			7.007
	SNL	7.400			7.400
	Other				0.000
Input/Output, Storage, and Networking		15.786	0.000	0.000	15.786
	LLNL	8.950			8.950
	LANL	2.336			2.336
	SNL	4.500			4.500
	Other				0.000
Post-processing Environments		10.605	0.000	0.000	10.605
	LLNL	2.195			2.195
	LANL	4.710			4.710
	SNL	3.700			3.700
	Other				0.000
Next-Generation Computing Technologies		0.000	0.000	0.000	0.000
	LLNL				0.000
	LANL				0.000
	SNL				0.000
	Other				0.000

Product / Project WBS	Lab	PEOPLE (\$Ms)	INFOSTRUCTURE (\$Ms)	CONTRACTS (\$Ms)	Total (\$Ms)
1.5.5 Facility Operations and User Support					
User Support Services		10.109	0.000	1.083	11.192
	LLNL	3.680		1.083	4.763
	LANL	4.229			4.229
	SNL	2.200			2.200
	Other				0.000
Collaborations		4.325	0.000	13.959	18.284
	LLNL	1.580		0.050	1.630
	LANL	1.140		0.215	1.355
	SNL	0.580		1.287	1.867
	Other	1.025		12.407	13.432
System and Environment Administration and Operations		31.951	43.800	0.000	75.751
	LLNL	12.400	30.000		42.400
	LANL	12.851	10.900		23.751
	SNL	6.700	2.900		9.600
	Other				0.000
Facilities, Network and Power		9.000	36.100	1.450	46.550
	LLNL	4.500	10.000		14.500
	LANL	3.000	20.600		23.600
	SNL	1.500	5.500	1.200	8.200
	Other			0.250	0.250
Common Computing Environment		9.615	0.000	0.795	10.410
	LLNL	3.673		0.795	4.468
	LANL	3.205			3.205
	SNL	2.737			2.737
	Other				0.000

	PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
Integrated Codes	144.907	0.000	20.500	165.407
Physics and Engineering Models	69.932	0.000	0.000	69.932
Verification and Validation	57.232	0.000	0.000	57.232
Computational Systems and Software Engineering	53.072	115.115	6.000	174.187
Facility Operations and User Support	65.000	79.900	17.287	162.187
Total	390.143	195.015	43.787	628.945
	62%	31%	7%	628.945

Month	Monthly Cost	Cumulative Cost
Oct	37.109	37.109
Nov	47.292	84.401
Dec	43.083	127.484
Jan	46.151	173.635
Feb	64.009	237.644
Mar	44.824	282.468
Apr	51.755	334.223
May	58.082	392.305
Jun	49.621	441.926
Jul	53.620	495.546
Aug	61.850	557.396
Sep	73.133	630.529

Appendix A. Glossary

2D	Two Dimensional
3D	Three Dimensional
ACE	ASC Computing Environment
ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AMR	Adaptive Mesh Refinement
API	Application Programming Interface
ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
ASCR	Advanced Scientific Computing Research (DOE/SC)
CBTF	Component-Based Tool Framework
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CRASH	Center for Radiative Shock Hydrodynamics
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
D&E	Development and Engineering
DDS	Distributed Data Services
DOE	Department of Energy
DSL	Domain-Specific Language
DSW	Directed Stockpile Work
FOUS	Facility Operations and User Support (WBS 1.5.5)
GA	General Availability
GB	Gigabytes
GPFS	Global Parallel File System
GPGPU	General-Purpose Computing on Graphics Processing Units
HMC	Hybrid Memory Cube

HPC	High Performance Computing
HPSS	High-Performance Storage System
HQ	ASC Headquarters
HTGL	High Temperature Gasdynamics Laboratory
I/O	Input/Output
IC	Integrated Codes
ID	Initial Delivery
IHPC	Inter-Site High Performance Computing
ISA	Tri-Labs Interconnection Security Agreement
ITIL	Information Technology Infrastructure Library
LA	Limited Availability
LAN	Local Area Network
LC	Livermore Computing Center
LDCC	Laboratory Data Communications Center
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
LVOC	Livermore Valley Open Campus
LWK	Lightweight Kernel
MOU	Memorandum of Understanding
MPI	Message Passing Interface
NAS	Network-Attached Storage
NFS	Network File System
NIC	Network Interface Chip
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NSA	National Security Agency
NSCC	National Security Computing Center
NSSB	National Security Sciences Building
NUMA	Non-Uniform Memory Access
nWBS	National Work Breakdown Structure
O SS	Open SpeedShop
OCF	Open Computing Facility

OS	Operating System
PCF	Predictive Capability Framework
PECOS	Predictive Engineering and Computational Sciences
PEM	Physics and Engineering Models
PLFS	Parallel Log Structured File System
PRISM	Prediction of Reliability, Integrity, and Survivability of Microsystems (Center for...)
PSAAP	Predictive Science Academic Alliance Program
PSI	Parallel Storage Interface
PSP	Predictive Science Panel
PUE	Power Usage Effectiveness
QDR	Quad Data Rate
QMU	Quantification of Margins and Uncertainties
QPX	Quad Floating Point
R&D	Research and Development
RAIT	Redundant Array of Independent Tapes
RANS	Reynolds-Averaged Navier Stokes
RHEL	Red Hat Enterprise Linux
SAN	Storage Area Network
SAP	Scalable Applications Preparations
SARAPE	Synchronized Account Request Automated Process
SC	Department of Energy's Office of Science
SCC	Strategic Computing Complex
SCR	Scalable/Check-Point Restart Code
SFI	Significant Finding Investigation
SIO	Scalable Input/Output
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SPEC	Science Partnership for Extreme-Scale Computing
SQA	Software Quality Assurance
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit

STAT	Stack Trace Analysis Tool
SU	Scalable Unit(s)
TBI	Thermonuclear Burn Initiative
TLCC	Tri-Lab Linux Capacity Cluster
TLS	Thread Level Speculation
TM	Transactional Memory
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
UQ	Uncertainty Quantification
V&V	Verification and Validation
VM	Virtual Machine
VTK	Visualization Toolkit
WAN	Wide Area Network
WBS	Work Breakdown Structure
WC Tool	Workload Characterization Tool

Appendix C. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.1	Capability Systems	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.2	Capacity Systems	Matt Leininger, LLNL, 925-422-4110, matt@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.3	Advanced Systems	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.4	System Software and Tools	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Sue Kelly, SNL, 505-845-9770, smkelly@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov Ron Oldfield, SNL, 505-284-9153, raoldfi@sandia.gov
1.5.4.6	Post-Processing Environments	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov David Daniel, LANL, 505-665-0883, ddd@lanl.gov David Rogers, SNL, 505-844-5323, dhroger@sandia.gov
1.5.4.7	Common Computing Environment	Moved
1.5.5	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Tom Klitsner, SNL, 505-844-1901, tklitsn@sandia.gov

WBS	Title	Contact
1.5.5.1	Facilities, Operations, and Communications	Retired
1.5.5.2	User Support Services	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.5.3	Collaborations	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Justine Johannes, SNL, 505-844-1994, jejohan@sandia.gov
1.5.5.4	System and Environment Administration and Operations	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.5.5.5	Facilities, Network, and Power	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov
1.5.5.6	Common Computing Environment	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Judy Sturtevant, SNL, 505-845-9448, jesturt@sandia.gov
1.5.1.4-TRI-001	Caltech, Center for Simulating Dynamic Response of Materials	Michael Ortiz, 626-395-4530, Ortiz@aeor.caltech.edu Mark Stalzer, 626-395-2521, stalzer@caltech.edu Susan Powell, 626-395-2909, spowell@cacr.caltech.edu
1.5.1.4-TRI-002	Purdue University, Center for Prediction of Reliability, Integrity and Survivability of Microsystems	Jayathi Murthy, 765-494-5701, jmurthy@ecn.purdue.edu Matt Potrawski, 765-496-9301, mpotraws@purdue.edu
1.5.1.4-TRI-003	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, moin@stanford.edu Gianluca Iaccarino, 650-723-9599, iops@stanford.edu Deb Michael, 650-725-2077, debmich@stanford.edu

WBS	Title	Contact
1.5.1.4- TRI-004	University of Michigan, Center for Radiative Shock Hydrodynamics	Paul Drake, 734-763-4072, rpdrake@umich.edu James Holloway, 734-936-3126, hagar@umich.edu
1.5.1.4- TRI-005	University of Texas, at Austin, Center for Predictive Engineering and Computational Sciences	Bob Moser, 512-471-0093, rmoser@mail.utexas.edu Chris Simmons, 512-232-2881, csim@colab.ices.utexas.edu

Appendix D.

WBS 1.5.1.4-TRI-001 Academic Alliance Centers

The Academic Alliance Centers project includes research activities at the funded academic centers as part of the PSAAP program, as listed below.

- California Institute of Technology (Caltech)
- Purdue University
- Stanford University
- University of Michigan
- University of Texas, at Austin

California Institute of Technology

The Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials

Caltech's Multidiscipline Simulation Center's overarching objective is the development of a multidisciplinary predictive science methodology focusing on high-energy-density dynamic response of materials and the demonstration of the methodology by means of a concerted and highly integrated experimental, computational, and analytical effort focusing on an overarching ASC-class problem: hypervelocity normal and oblique impact of projectiles on metallic and non-metallic targets, at velocities up to 10 km/sec. Hypervelocity impact gives rise to pressures in the Mbar range and strain-rates up to 10^{11} /sec., providing a grand-challenge problem in predictive science that is also well-matched to the direct interests of the NNSA mission.

Planned activities in FY12:

The main Year 4 overarching activities will be:

- Lagrangian Optimal Transportation Meshfree and Eulerian simulations of oblique hypervelocity impact of single multiple aluminum plates and nylon projectiles in the 5–10 Km/sec. impact velocity range, including verification, performance analysis, and UQ runs using the UQ pipeline and multiple performance measures.
- Hypervelocity tests on multiple aluminum target plates and nylon projectiles in the 5–10 Km/sec. range, including experimental setup and UQ tests using profilometry, velocimetry, shadowgraph and two-dimensional (2D) fragment-capture measurements.

Expected deliverables in FY12:

- Complete legacy-UQ analysis of the nylon/aluminum hypervelocity experiment including multiple performance measures
- Parallel implementation of Optimal Transportation Meshfree with dynamic load balancing
- Sustained petascale UQ heterogeneous/distributed calculations using the UQ pipeline
- Validation of the multi-scale strength model of iron using laser-driven Rayleigh-Taylor experiments
- UQ analysis of the multi-scale strength and fracture models of iron
- Formulation of “plasma-aware” material models
- Calculation of emission spectra during the impact flash in hypervelocity experiments

Preliminary planned activities in FY13:

- Nylon/iron/aluminum hypervelocity Eulerian (VTF) and Lagrangian (Optimal Transportation Meshfree) simulations for UQ analysis.
- Nylon/iron/aluminum hypervelocity tests for UQ analysis with profilometry, velocimetry, shadowgraph, spectrometric, and 3D fragment-capture measurements.
- Final deployment and release of the UQ pipeline as a heterogeneous and distributed simulation and optimization framework, including Eulerian (VTF) and Lagrangian (Optimal Transportation Meshfree) solvers and verification/validation suite.
- Final deployment and release of optimal legacy-data UQ protocol within the UQ pipeline.
- Final deployment and release of library of multi-scale, phase, and plasma-aware strength, equation of state, and fracture material models including verification/validation suite.
- Final archiving and release of in-house material and hypervelocity experimental data.

Purdue***Center for Prediction of Reliability, Integrity, and Survivability of Microsystems***

The overall objective of PRISM is to accelerate substantially the integration of microelectromechanical systems technologies into civilian and defense systems. PRISM aims to significantly improve understanding of the long-term reliability of microelectromechanical systems and survivability in harsh environments by simulating rigorously, and at multiple scales, the physics of failure, accounting for the coupled electrical, mechanical, thermal, and materials behavior of microelectromechanical systems, from atoms to devices. Advanced simulation software developed by the Center will be encapsulated in an integrated simulation system, MEMOSA.

Planned activities in FY12

- Perform experiments on multi-cycle contact, pull-in, pull-out, and creep using the PRISM Generation 6 switch
- Refine metal-dielectric contact model to include plasticity; V&V
- Refine creep model by calibration against new Generation 6 PRISM switch creep data, and validation of MEMOSA simulations against measurements
- Refine multiple trap depth dielectric charging model using Generation 6 dielectric charging measurements on PRISM switch; V&V
- Integrate the refined physics models in MEMOSA; V&V.
- Integrate ES-BGK rarefied gas dynamics model with immersed boundary method in MEMOSA and V&V of the model
- Integrate simulations of 1) multi-cycle contact simulations of PRISM switch with coupled fluid-structure-electrostatics interactions, contact, and dielectric charging, and prediction of dielectric charging rate; 2) sustained long-term contact simulations in the presence of dielectric charging, and prediction of pull-out voltage after sustained contact; and 3) long-term creep simulations and prediction of gap-versus-time

Expected deliverables in FY12:

- Integrated simulations of multi-cycle contact simulations of PRISM switch with coupled fluid-structure-electrostatics interactions, contact and dielectric charging, and prediction of dielectric charging rate with UQ
- Sustained long-term contact simulations in the presence of dielectric charging, and prediction of pull-out voltage after sustained contact with UQ
- Long-term creep simulations and prediction of gap-versus-time under actuation at different voltages with UQ

Preliminary planned activities in FY13:

- Continue experiments on multi-cycle contact, pull-in, pull-out, and creep using the PRISM switch
- Refine integrated simulations of multi-cycle contact simulations of PRISM switch with coupled fluid-structure-electrostatics interactions, contact and dielectric charging, and prediction of dielectric charging rate with UQ
- Refine sustained long-term contact simulations in the presence of dielectric charging, and prediction of pull-out voltage after sustained contact with UQ
- Refine long-term creep simulations and prediction of gap-versus-time under actuation at different voltages with UQ

Stanford University

The Center for Predictive Simulations of Multi-Physics Flow Phenomena with Application to Integrated Hypersonic Systems

The objective of the Center is to characterize the operability limits of an air-breathing hypersonic vehicle and associated propulsion system using predictive multi-physics simulations. The primary focus will be on the “unstart” failure mode triggered by thermal choking. Air-breathing hypersonic vehicles are envisioned as a means for reliable low-cost access to space. These vehicles are highly integrated systems whose performance depends on complex physics and the interactions between all of its components. Current state-of-the-art simulation capabilities cannot predict these systems reliably, particularly near their operability limits.

World-class experimental facilities in Stanford’s High Temperature Gasdynamics Laboratory (HTGL) will be used to conduct tightly integrated validation experiments for the key component physics and models. The development and implementation of UQ methods for very large systems will be an integral part of the effort at both the component and the system levels. Novel verification methods for high-fidelity simulations will also be developed and implemented. Stanford will leverage advanced computer science methods developed at Stanford to directly impact simulation tools and ensure scalability, program correctness, and portability to future platforms with very large numbers of cores.

Planned activities in FY12:

- Perform multiple-gate QMU analysis of full system, including aleatoric and epistemic uncertainties
- Apply novel balance of error and uncertainties approach to full system reacting flow simulations
- Integrate adjoint solver for estimation of sensitivities and numerical errors into simulations at all levels (full-system, sub-system, and unit problem)
- Implement full turbulent version of the Joe RANS solver into the DSL Liszt; run on multiple backends, including GPUs and heterogeneous systems
- Initiate experimental campaign of the novel model combustor at Stanford and carry out both RANS and LES simulations
- Perform Monte-Carlo experiment and corresponding simulations of shock/boundary layer interaction
- Investigate epistemic uncertainties in models of shock/boundary layer interaction and mixing using both experimental data and LES studies
- Initiate simulation activities for the HIFIRE (AirForce) scramjet vehicle

Expected deliverables in FY12:

- Fully integrated RANS and adjoint solver for supersonic reacting flows in complex geometries
- Epistemic uncertainty model for shock/boundary layer interactions

- Demonstrated capability to implement multi-physics solver in domain-specific language (DSL) Liszt, and demonstrated performance on the heterogeneous architectures expected of future exascale systems

Preliminary planned activities in FY13

- Multiple-gate QMU analysis of full system with fully integrated RANS/adjoint solver

University of Michigan

The Center for Radiative Shock Hydrodynamics (CRASH) is advancing predictive science in the nationally important area of radiation hydrodynamics via a unified, multi-prong approach. To substantially improve the ability to perform predictive simulations of high-energy-density and astrophysical flows, Center researchers:

- Continue to develop software for radiation hydrodynamics to serve as a test bed for development and V&V of radiation hydrodynamics modeling elements
- Continue to develop a system for extensive V&V of the software
- Extended an existing experimental effort, centered on radiative shocks, to quantify uncertainties in the experiments and obtain data focused on improving the predictive capability of the code
- Simulate this sequence of experiments and assess the predictive capability of the simulations
- Establish a doctoral program track for predictive science and engineering

Planned activities in FY12:

- Carry out the sequence of computational run sets described in the 2010 Annual Review and in the technical report. These run sets cover a number of issues including numerical parameters and variations in model fidelity.
- Using this sequence of run sets, the methodology, and data from the Year 1 through 3 CRASH experiments, develop predictive distributions and certainty of predictions for the Year 5 experiments.
- Carry out studies using BATSRUS/CRASH and PDT to assess the errors associated with using multi-group diffusion in CRASH for the radiation energy transport, including the studies suggested by the 2010 review. Seek to develop source terms that might compensate for these errors in the diffusion model.
- Fully activate the laser package and release CRASH v 3.0.
- Define and carry out the PY4 experiment consistent with the recommendations of the 2009 and 2010 reviews, to observe the effects of symmetric nozzles and obtain initial data from the Year 5 experimental configuration.

Expected deliverables in FY12:

- A report, as part of the annual technical report, reporting the results of the computational run sets, the subsequent analysis, and the resulting predictive

distributions; of the quantitative evaluation of errors in diffusion methods; of new features of the CRASH code and the new V&V results; and of experiments

Preliminary planned activities in FY13:

- Perform a sequence of run sets to be able to develop predictive probability distributions for the year 5 experiments
- Using the output of these run sets, develop these predictive probability distributions
- Support evolution of the code as needed for these efforts and continued broadening of validation studies and applications
- Continue to analyze and report data from PY1 through PY4 experiments and prepare for PY 5 experiments

University of Texas

The Center for Predictive Engineering and Computational Sciences

The goal of the Predictive Engineering and Computational Sciences (PECOS) Center is to develop next generation advanced computational and UQ methods for predictive simulation of multi-scale, multi-physics phenomena relevant to the NNSA, and to apply these methods to the problem of reentry of vehicles into the atmosphere.

Simulation of vehicle reentry into the atmosphere requires modeling of the interaction of extremely high temperature gas flows with the high temperature response of materials, in particular the vehicle's thermal protection system. The high gas temperatures produce chemical dissociation, thermal non-equilibrium, and ionization. Radiative heat transfer is an important part of the heat load on the vehicle, while transition and turbulence greatly enhance the rate of heat transfer. During reentry, the ablative thermal protection system responds via pyrolysis, chemical reaction, and formation and mechanical degradation of a refractory char layer. Models of these high-energy, multi-scale, multi-physics phenomena have been integrated into a unified simulation code, designed to support predictive simulation.

Planned activities in FY12:

- Integrate FIN-S with updated physics models in radiation, turbulence, transport, and chemistry as guided by validation
- Continue simulation of high-speed boundary layers via DNS
- Continue development of efficient variants of the Stochastic Newton method
- Explore fault-tolerance implementations for statistical calibration problems
- Explore adjoint-enhanced forward propagation of uncertainty

Expected deliverables in FY12:

- Completion of forward propagation of uncertainty in the full system simulation using calibrated single physics and calibrated multi-physics

- Completion of calibration of coupled ablation/chemistry with ablation boundary layer DNS data
- Completion of calibration of coupled chemistry/turbulence models using DNS data
- Addition of coupled multi-physics manufactured solutions across all modeling domains

Preliminary planned activities in FY13:

- Complete forward propagation of uncertainty in the full system simulation using calibrated and validated models across all domains
- Develop additional adjoint-enabled modeling codes as guided by validation processes
- Continue working on validation methodology using the information-theoretic interpretation of Bayes theorem
- Complete new data acquisition in surface nitridation/oxidation using optimal experimental design

Appendix E. ASC Obligation/Spend Plan




Figure E-1. ASC obligation/spend plan for FY12.

